Cisco Systems

Send documentation comments to mdsfeedback-doc@cisco.com

Cisco MDS 9000 Family Fabric Manager Configuration Guide

Cisco MDS SAN-OS Release 3.0(1) Through 3.0(2)
Cisco MDS 9000 FabricWare Release 2.x
May 2006

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-8007-02

Cisco MDS 9000 Family Fabric Manager Configuration Guide

New and Changed Information

This document provides release-specific information for each new and changed feature in Cisco Fabric Manager and Cisco MDS SAN-OS Release 3.x software. The *Cisco MDS 9000 Family Fabric Manager Configuration Guide* is updated to address each new and changed feature in the Cisco MDS SAN-OS Release 3.x software. The latest version of this document is available at the following website: http://www.cisco.com/univercd/cc/td/doc/product/sn5000/mds9000/index.htm

Tip The configuration guides created for earlier releases are also listed at the website. Each guide addresses the features introduced or available in those releases. Select and view the configuration guide that applies to the software installed in your switch.

To check for additional information about this release, refer to the *Cisco MDS 9000 Family Release Notes* available at the following Cisco Systems website: http://www.cisco.com/en/US/products/ps5989/prod_release_notes_list.html

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

Table 1 New and Changed Features for Release 3.x

Feature	GUI Change	Description	Changed in Release	Where Documented
ESCON Style Port Configuration Display	ESCON Style Port checkbox added to the Port Configuration Table	Allows the user to see which ports are available or prohibited in Device Manager.	3.0(2)	Chapter 31, "Configuring FICON"
FICON Configuration Locking	A confirmation dialog box displayed if the FICON VSAN information is changed.	Allows the user to save any changes to the FICON configuration since the last refresh.	3.0(2)	Chapter 31, "Configuring FICON"
Highlighting the E/TE ports or multiple interfaces	Added a tooltip	In Device Manager, the FICON port configuration table when clicked the port address index column is highlighted red which identifies the E/TE ports or multiple interfaces.	3.0(2)	Chapter 31, "Configuring FICON"

Table 1 New and Changed Features for Release 3.x (continued)

Feature	GUI Change	Description	Changed in Release	Where Documented
FICON Port Configuration for Multiple VSANs	Added Port Configuration and Files tabs to the FICON table	The Port Configuration tab allows the user to block the required ports. The Files tab allows the user to open and view any of the configuration files.	3.0(2)	Chapter 31, "Configuring FICON"
Requiring privacy password for user creation	None	You must be logged into Fabric Manager or Device Manager with your password and privacy password to create users.	3.0(1)	Chapter 33, "Configuring Users and Common Roles"
SNMP over TCP/IP	None—updated Server Properties file	Allows SNMP messages to be transported over TCP rather than UDP for management traffic on the out-of-band Ethernet management port (mgmt0).	3.0(1)	Chapter 3, "Fabric Manager Server"
EMC Call Home	None—updated Server Properties file	Allows the forwarding of traps as XML data using email, according to EMC specifications	3.0(1)	Chapter 3, "Fabric Manager Server"
Syslog messages	Syslog link	Click a syslog link to directly access the syslog files	3.0(1)	Chapter 3, "Fabric Manager Server"
Traffic statistics	Hot links for traffic data grouped by Host	If the data is grouped by Hosts, you can click on the host link to get traffic statistics	3.0(1)	Chapter 3, "Fabric Manager Server"
Reporting of enclosure aliases	Report Aliases Preference Added	Checking this check box in Preferences > General shows the aliases for the enclosures when you mouse over them in the Fabric pane.	3.0(1)	Chapter 4, "Fabric Manager Client"
Fabric Manager Install Wizard	Install Wizard updated to support Motorola-based clients	Supports installation on Motorola-based systems such as the Apple Power PC.	3.0(1)	Chapter 4, "Fabric Manager Client"
Filtering by user defined groups	Group menu item added to switch and host context menus	Allows the definition of custom groups containing switches, or hosts and storage devices, to enable the filtering of information that is not relevant to that group. Filtering applies to Cisco Fabric Manger topology map, information tables (switch parameters), and Fabric Manager Server (FMS) reports.	3.0(1)	Chapter 4, "Fabric Manager Client"

Table 1 New and Changed Features for Release 3.x (continued)

Feature	GUI Change	Description	Changed in Release	Where Documented
Filtering by user defined groups	Groups folder added	Allows the definition of custom groups containing switches, or hosts and storage devices, to enable the filtering of information that is not relevant to that group. Filtering applies to Cisco Fabric Manger topology map, information tables (switch parameters), and Fabric Manager Server (FMS) reports.	3.0(1)	Chapter 4, "Fabric Manager Client"
Fabric pane layout saving	Save Layout menu item added to general context menu	This feature, which was removed in Release 2.x, has been added to Release 3.0.	3.0(1)	Chapter 4, "Fabric Manager Client"
Export Device Manager image	Export Image menu item on Device menu	Saves a snapshot of the Device Manager display to the local PC.	3.0(1)	Chapter 5, "Device Manager"
Print Device Manager image	Print menu item on Device menu	Prints a snapshot of the Device Manager display to the local PC	3.0(1)	Chapter 5, "Device Manager"
Browser preference for Device Manager display	Browser Preferences setting in Preferences dialog box	Allows you to set the type of browser you use to display the Device Manager (Netscape or Mozilla-based)	3.0(1)	Chapter 5, "Device Manager"
Change in Device Manager Inventory display.	Module ID column on Inventory display	Displays the module ID.	3.0(1)	Chapter 5, "Device Manager"
Performance Collections	Performance Manager Collection Wizard Moved to Web Services	Collections are now configured using the Fabric Manager Web Services client under the Admin > Configuration > Collections screen.	3.0(1)	Chapter 6, "Fabric Manager Web Services"
Searchable statistical information in the Fabric Manager Web Services client	Searchability icon displayed in certain table columns	Click on the search icon in some tables to search for items listed in that column.	3.0(1)	Chapter 6, "Fabric Manager Web Services"
Event Forwarding	New Event screen	Events logged by Cisco FMS can be forwarded to users via e-mail, or as SNMP Traps to network management systems.	3.0(1)	Chapter 6, "Fabric Manager Web Services"
Data collection auto update	Data Collection tab	Provides an automatic way to keep data collection definitions current.	3.0(1)	Chapter 6, "Fabric Manager Web Services"
SCSI flow statistics reports	SCSI flow tab	Provides a real-time view of SCSI flow statistics, which include logical unit number (LUN) level throughput, I/Os per second, response times, and SCSI error information.	3.0(1)	Chapter 6, "Fabric Manager Web Services"

Table 1 New and Changed Features for Release 3.x (continued)

Feature	GUI Change	Description	Changed in Release	Where Documented
Server performance summary reports	Server performance tab	Provides a view of summary throughput, errors, and discards statistics for all connections on paths from a server to its storage devices.	3.0(1)	Chapter 6, "Fabric Manager Web Services"
Custom report performance charts	Custom report tab	Allows throughput performance charts to be optionally embedded for each table entry.	3.0(1)	Chapter 6, "Fabric Manager Web Services"
Configurable RRD	Database menu option on Admin tab	Allows configuration of the number of samples saved for each resolution.	3.0(1)	Chapter 6, "Fabric Manager Web Services"
Fabric Manager Web Services administration	Reorganized Admin area	Reorganization of certain areas of the Admin tab, such as Performance and Fabrics	3.0(1)	Chapter 6, "Fabric Manager Web Services"
Communities configuration	New Communities screen	Allows you to configure communities.	3.0(1)	Chapter 6, "Fabric Manager Web Services"
Local Roles configuration	New Web Local Roles screen	Allows you to configure local roles.	3.0(1)	Chapter 6, "Fabric Manager Web Services"
Performance Prediction	New Performance Prediction tab	Provides a way to more reliably predict when storage network connections will become overutilized.	3.0(1)	Chapter 6, "Fabric Manager Web Services"
Performance statistics	New Last Updated column on performance statistics	Displays the date and time the statistics were last updated.	3.0(1)	Chapter 6, "Fabric Manager Web Services"
Crossbar graceful shutdown	Out of Service menu item on Supervisor context menu	Provides procedures for gracefully shutting down the crossbars on the Cisco MDS 9500 Series Directors.	3.0(1)	Chapter 16, "Managing System Hardware"
Crossbar removal procedures	---	Provides the recommended procedures to prepare to remove crossbars from Cisco MDS 9500 Series Directors.	3.0(1)	Chapter 16, "Managing System Hardware"
Supervisor-2 module support	---	Includes support for the following Supervisor-2 module features: • Configuring modem parameters on the console port and COM1 port. • Allowing 1000 Mbps speed on the management port.	3.0(1)	Chapter 11, "Initial Configuration"
Generation 2 switching module support	---	Describes how to configure interfaces on the Generation 2 Fibre Channel switching modules.	3.0(1)	Chapter 19, "Configuring Generation 2 Switching Modules"
CFS support for allowed domain ID lists	Allowed DomainIds tab under VSAN, Domain Manager	Allows the allowed domain ID lists to be distributed in the fabric using the CFS infrastructure.	3.0(1)	Chapter 22, "Configuring Domain Parameters"

Table 1 New and Changed Features for Release 3.x (continued)

Feature	GUI Change	Description	Changed in Release	Where Documented
CFS over IP	---	Allows CFS distributions over IP connections.	3.0(1)	Chapter 12, "Using the CFS Infrastructure"
Supervisor module management procedures	---	Includes the following recommended supervisor module management procedures: • Preparing to remove supervisor modules from Cisco MDS 9500 Series Directors containing both Generation 1 and Generation 2 switching modules. • Migrating from Supervisor 1 modules to Supervisor 2 modules in the Cisco MDS 9500 Series Directors.	3.0(1)	Chapter 13, "Software Images"
IVR zone rename	Rename menu item on Edit menu in Edit Local Full Zone Database dialog box	Allows the renaming of IVR zones	3.0(1)	Chapter 25, "Configuring Inter-VSAN Routing"
IVR zone set rename	Rename menu item on Edit menu in Edit Local Full Zone Database dialog box	Allows the renaming of IVR zone sets	3.0(1)	Chapter 25, "Configuring Inter-VSAN Routing"
Active IVR zone set (IVZS) copy	Copy menu item on Edit menu in Edit Local Full Zone Database dialog box	Allows copying the active IVZS to the full IVZS to be edited and reactivated.	3.0(1)	Chapter 25, "Configuring Inter-VSAN Routing"
Active IVR topology copy	Discrepancies tab and Domains tab visible after clicking IVR, CFS tab	Allows copying the active IVR topology to the manually configured IVR topology	3.0(1)	Chapter 25, "Configuring Inter-VSAN Routing"
Increased zone limit per VSAN	None	Increases the maximum number of zones per VSAN from 2000 to 8000.	3.0(1)	Chapter 26, "Configuring and Managing Zones"
In-order-delivery enhancement	---	Ensures that frames are delivered in order within the switch latency drop period.	3.0(1)	Chapter 40, "Configuring FC-SP and DHCHAP"
CFS support for RCSN	CFS tab under VSAN, Domain Manager, Advanced	Allows the RSCN timer value to be distributed in the fabric using the CFS infrastructure.	3.0(1)	Chapter 29, "Managing FLOGI, Name Server, FDMI, and RSCN Databases"
RSCN timer value configuration	RSCN Event tab under VSAN, Domain Manager, Advanced	Allows the RSCN timer value to be configured.	3.0(1)	Chapter 29, "Managing FLOGI, Name Server, FDMI, and RSCN Databases"

Table 1 New and Changed Features for Release 3.x (continued)

Feature	GUI Change	Description	Changed in Release	Where Documented
FICON port numbering	None	Provides information on the changed default port numbering scheme for Generation 2 hardware and how to assign port numbers when a switch has more than 255 ports.	3.0(1)	Chapter 31, "Configuring FICON"
McDATA native interoperability	McDATA tab	Describes how to configure McDATA native mode interoperability.	3.0(1)	Chapter 32, "Advanced Features and Concepts"
Certificate authorities and digital certificates	PKI item under Security	Describes interoperating with Certificate Authorities and using digital certificates for secure communication with peers.	3.0(1)	Chapter 38, "Configuring Certificate Authorities and Digital Certificates"
IKE identity	IKE item under Security	Allows an IKE identity host name or IP address to be specified in the IPsec domain.	3.0(1)	Chapter 39, "Configuring IPsec Network Security"
IKE key host name	Pre-Shared AuthKey tab under Security, IKE	Allows IKE identity host name to be specified instead of an IP address for preshared keys.	3.0(1)	Chapter 39, "Configuring IPsec Network Security"
AAA server enhancements	AAA folder under Security	Includes the following AAA server enhancements: • Validating the availability of remote AAA servers. • Allowing users to specify a remote AAA server name at login. Displaying AAA server statistics.	3.0(1)	Chapter 35, "Configuring RADIUS and TACACS+"
MS CHAP	AuthTypeMSCHAP check box under Security, AAA, General tab	Provides support for the Microsoft Challenge Handshake Authentication Protocol (MSCHAP).	3.0(1)	Chapter 35, "Configuring RADIUS and TACACS+"
Fabric binding for Fibre Channel	Fabric Binding under VSAN, Domain Manager	Supports fabric binding for Fibre Channel VSANs as well as FICON VSANs.	3.0(1)	Chapter 42, "Configuring Fabric Binding"
FCIP tape read acceleration	---	Supports tape read acceleration over FCIP interfaces as well as tape write acceleration.	3.091)	Chapter 43, "Configuring FCIP"
iSCSI server load balancing (iSLB)	---	Provides information about how to easily configure large iSCSI deployments.	3.0(1)	Chapter 45, "Configuring iSCSI"
iSNS cloud discovery	---	Provides information to iSNS on the various interfaces reachable from an initiator by partitioning the interfaces on a switch into disjoint IP clouds.	3.0(1)	Chapter 45, "Configuring iSCSI"

Table 1 New and Changed Features for Release 3.x (continued)

Feature	GUI Change	Description	Changed in Release	Where Documented
IPv6 access control lists (IPv6-ACLs)	IP ACL Wizard on the Tools menu	Describes the support for IPv6-ACLs.	3.0(1)	Chapter 48, "Configuring IPv6 for Gigabit Ethernet Interfaces"
Dynamic initiator modes	---	Allows configuration of dynamic initiator modes iSCSI, iSLB, and deny log in to the MDS switch.	3.0(1)	Chapter 45, "Configuring iSCSI"
IPv6	None	Provides support for IP version 6 (IPv6).	3.0(1)	Chapter 46, "Configuring IP Services" Chapter 48, "Configuring IPv6 for Gigabit Ethernet Interfaces"
Call Home enhancement	---	Allows customization of alert group messages.	3.0(1)	Chapter 58, "Configuring Call Home"
SAN extension tuner enhancement	---	Describes how to assign tape read and write commands to N ports.	3.0(1)	Chapter 38, "Configuring the SAN Extension Tuner"
QoS behavior	---	Provides information about the behavior of QoS with different combinations of Generation 1 and Generation 2 switching modules.	3.0(1)	Chapter 60, "Configuring Fabric Congestion Control and QoS"

Preface

This preface describes the audience, organization, and conventions of the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*. It also provides information on how to obtain related documentation.

Audience

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

Organization

This guide is organized as follows:

Chapter	Title	Description
Chapter 1	Overview	Presents an overview of the Cisco MDS 9000 Family of multilayer switches and directors.
Chapter 2	Installing Cisco MDS SAN-OS and Fabric Manager	Provides a brief overview of Fabric Manager components and capabilities, and information on installation and launching the applications.
Chapter 3	Fabric Manager Server	Provides in-depth descriptions of GUI and capabilities for the Fabric Manager Server.
Chapter 4	Fabric Manager Client	Provides in-depth descriptions of GUI and capabilities for the Fabric Manager.
Chapter 5	Device Manager	Provides in-depth descriptions of GUI and capabilities for the Device Manager.
Chapter 6	Fabric Manager Web Services	Provides in-depth descriptions of GUI and capabilities for the Fabric Manager Web Client.
Chapter 7	Performance Manager	Provides overview of Performance Manager architecture.
Chapter 8	Authentication in Fabric Manager	Describes the authentication schemes between Fabric Manager components and fabric switches.

Chapter	Title	Description
Chapter 9	Cisco Traffic Analyzer	Describes installing and launching Cisco Traffic Analyzer from Performance Manager.
Chapter 10	Obtaining and Installing Licenses	Describes license types, procedure, installation, and management for the Cisco MDS SAN-OS software.
Chapter 11	Initial Configuration	Provides initial switch configuration options and switch access information.
Chapter 12	Using the CFS Infrastructure	Explains the use of the Cisco Fabric Services (CFS) infrastructure to enable efficient database distribution.
Chapter 13	Software Images	Describes how to install and upgrade software images
Chapter 14	Managing Configuration Files	Describes the initial configuration of the switches using the configuration files so they can be accessed by other devices
Chapter 15	Configuring High Availability	Describes the high availability feature including switchover mechanisms.
Chapter 16	Managing System Hardware	Explains switch hardware inventory, power usage, power supply, module temperature, fan and clock modules, and environment information.
Chapter 17	Managing Modules	Explains how to display and analyze the status of each module and specifies the power on and power off process for modules.
Chapter 18	Configuring Interfaces	Explains Generation 1 and Generation 2 module port and operational state concepts in Cisco MDS 9000 Family switches and provides details on configuring ports and interfaces.
Chapter 19	Configuring Generation 2 Switching Modules	Explains configuration concepts for Generation 2 module ports and interfaces.
Chapter 20	Configuring Trunking	Explains TE ports and trunking concepts.
Chapter 21	Configuring PortChannels	Explains PortChannels and load balancing concepts and provides details on configuring PortChannels, adding ports to PortChannels, and deleting ports from PortChannels.
Chapter 22	Configuring Domain Parameters	Explains the Fibre Channel domain (fcdomain) feature, which includes principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions.
Chapter 23	Configuring and Managing VSANs	Describes how virtual SANs (VSANs) work, explains the concept of default VSANs, isolated VSANs, VSAN IDs, and attributes, and provides details on how to create, delete, and view VSANs.

Chapter	Title	Description
Chapter 24	Creating Dynamic VSANs	Defines the Dynamic Port VSAN Membership (DPVM) feature that is used to maintain fabric topology when a host or storage device connection is moved between two Cisco MDS switches.
Chapter 25	Configuring Inter-VSAN Routing	Provides details on sharing resources across VSANs using the inter-VSAN Routing (IVR) feature.
Chapter 26	Configuring and Managing Zones	Defines various zoning concepts and provides details on configuring a zone set and zone management features.
Chapter 27	Distributing Device Alias Services	Describes the use of the Distributed Device Alias Services (device alias) to distribute device alias names on a fabric-wide basis.
Chapter 28	Configuring Fibre Channel Routing Services and Protocols	Provides details and configuration information on Fibre Channel routing services and protocols.
Chapter 29	Managing FLOGI, Name Server, FDMI, and RSCN Databases	Provides name server and fabric login details required to manage storage devices and display registered state change notification (RSCN) databases.
Chapter 30	Discovering SCSI Targets	Describes how the SCSI LUN discovery feature is started and displayed.
Chapter 31	Configuring FICON	Provides details on the FI-bre CON-nection (FICON) interface, fabric binding, and the Registered Link Incident Report (RLIR) capabilities in Cisco MDS switches.
Chapter 32	Advanced Features and Concepts	Describes the advanced configuration features—time out values, fctrace, fabric analyzer, world wide names, flat FC IDs, loop monitoring, and interoperating switches.
Chapter 33	Configuring Users and Common Roles	Describes how to configure users and common roles.
Chapter 34	Configuring SNMP	Provides details on how you can use SNMP to modify a role that was created using CLI.
Chapter 35	Configuring RADIUS and TACACS+	Discusses the AAA parameters, user profiles, and RADIUS authentication security options provided in all switches in the Cisco MDS 9000 Family and provides configuration information for these options.
Chapter 36	Configuring IPv4 Access Control Lists	Describes the IPv4 static routing feature and its use to route traffic between VSANs.
Chapter 48	Configuring IPv6 for Gigabit Ethernet Interfaces	Describes the IPv6 static routing feature and its use to route traffic between VSANs.

Chapter	Title	Description
Chapter 38	Configuring Certificate Authorities and Digital Certificates	Describes how to interoperate with Certificate Authorities (CAs) and use digital certificates for secure, scalable communication.
Chapter 39	Configuring IPsec Network Security	Provides details on the digital certificates, IP Security Protocol (IPsec) open standards, and the Internet Key Exchange (IKE) protocol that it uses to handle protocol and algorithm negotiation.
Chapter 40	Configuring FC-SP and DHCHAP	Describes the DHCHAP protocol, an FC-SP protocol, that provides authentication between Cisco MDS 9000 Family switches and other devices.
Chapter 41	Configuring Port Security	Provides details on port security features that can prevent unauthorized access to a switch port in the Cisco MDS 9000 Family.
Chapter 42	Configuring Fabric Binding	Describes the fabric binding security feature for VSANs, which ensures that ISLs are only enabled between specific switches.
Chapter 43	Configuring FCIP	Describes how the switch allows IP hosts to access Fibre Channel storage using the iSCSI protocol.
Chapter 44	Configuring the SAN Extension Tuner	Explains the SAN extension tuner (SET) feature that optimizes FCIP performance.
Chapter 45	Configuring iSCSI	Describes the iSCSI feature that is specific to the IPS module and is available in the Cisco MDS 9200 Switches or Cisco MDS 9500 Directors.
Chapter 46	Configuring IP Services	Provides details on IP over Fibre Channel (IPFC) services and provides configuring IPFC, virtual router, and DNS server configuration information.
Chapter 47	Configuring IP Storage	Provides details on extending the reach of Fibre Channel SANs by connecting separated SAN islands together through IP networks using FCIP, and allowing IP hosts to access FC storage using the iSCSI protocol.
Chapter 48	Configuring IPv6 for Gigabit Ethernet Interfaces	Describes the IPv6 protocol support provided by Cisco MDS 9000 Family switches.
Chapter 49	Configuring SCSI Flow Services and Statistics	Describes the SCSI flow services and SCSI flow statistics, the Intelligent Storage Services.
Chapter 50	Configuring Fibre Channel Write Acceleration	Describes Fibre Channel Write Acceleration support and configuration.
Chapter 51	Configuring SANTap	Describes SANTap support and configuration.
Chapter 52	Configuring NASB	Describes NASB support and configuration.

Chapter	Title	Description
Chapter 53	Network Monitoring	Describes how to use Fabric Manager monitoring features.
Chapter 54	Performance Monitoring	Provides details on using Performance Manager.
Chapter 55	Configuring RMON	Provides details on using RMONs to configure alarms and events.
Chapter 56	Monitoring Network Traffic Using SPAN	Describes the Switched Port Analyzer (SPAN), SPAN sources, filters, SPAN sessions, SD port characteristics, and configuration details.
Chapter 57	Configuring System Message Logging	Describes how system message logging is configured and displayed.
Chapter 58	Configuring Call Home	Provides details on the Call Home service and includes information on Call Home, event triggers, contact information, destination profiles, and e-mail options.
Chapter 59	Configuring Call Home	Describes how the fabric configuration server (FCS) feature is configured and displayed.
Chapter 60	Configuring Fabric Congestion Control and QoS	Provides details on the quality of service (QoS) and Fibre Channel Congestion Control (FCC) features.
Chapter 61	Configuring Port Tracking	Provides information about a port tracking feature that provides a faster recovery from link failures.
Chapter 62	Troubleshooting Your Fabric	Provides details on fctrace, fcping, and other troubleshooting tools.
Chapter 63	Management Software FAQ	Provides details on troubleshooting Fabric Manager.
Chapter 64	Monitoring System Processes and Logs	Provides information on displaying system processes and status. It also provides information on configuring core and log files, HA policy, heartbeat and watchdog checks, and upgrade resets.
Appendix A	Cisco Fabric Manager Unsupported Feature List	Provides a table of procedures, organized by best performed by the CLI, Fabric Manager, or Device Manager.
Appendix B	Interface Nonoperational Reason Codes	Explains the reason codes for why an interface is operationally down.
Appendix C	Managing Cisco FabricWare	Provides information on managing products running Cisco FabricWare.

Document Conventions

Command descriptions use these conventions:

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

Screen examples use these conventions:

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

This document uses the following conventions:

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

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

Related Documentation

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

Release Notes

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

Compatibility Information

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

Regulatory Compliance and Safety Information

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

Hardware Installation

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

Send documentation comments to mdsfeedback-doc@cisco.com

Cisco Fabric Manager

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

Command-Line Interface

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

Troubleshooting and Reference

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

Installation and Configuration Note

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

Obtaining Documentation

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

Send documentation comments to mdsfeedback-doc@cisco.com

Cisco.com

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

http://www.cisco.com/techsupport

You can access the Cisco website at this URL:

http://www.cisco.com

You can access international Cisco websites at this URL:

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

Product Documentation DVD

The Product Documentation DVD is a comprehensive library of technical product documentation on a portable medium. The DVD enables you to access multiple versions of installation, configuration, and command guides for Cisco hardware and software products. With the DVD, you have access to the same HTML 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= or DOC-DOCDVD=SUB) from Cisco Marketplace at this URL:

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

Ordering Documentation

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 submit comments about Cisco documentation 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 will find information about how to:

- 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, security notices, and security responses for Cisco products is available at this URL:

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

To see security advisories, security notices, and security responses as they are updated in real time, you can subscribe to the Product Security Incident Response Team Really Simple Syndication (PSIRT RSS) feed. Information about how to subscribe to the PSIRT RSS feed is found at 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 have identified a vulnerability in a Cisco product, contact PSIRT:

- For Emergencies only—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.
- For 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 (for example, GnuPG) to encrypt any sensitive information that you send to Cisco. PSIRT can work with information that has been encrypted with PGP versions 2.*x* through 9.*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.

If you do not have or use PGP, contact PSIRT at the aforementioned e-mail addresses or phone numbers before sending any sensitive material to find other means of encrypting the data.

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)—An existing 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 operations 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 the network is impaired, while 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.

- The *Cisco Product Quick Reference Guide* is a handy, compact reference tool that includes brief product overviews, key features, sample part numbers, and abbreviated technical specifications for many Cisco products that are sold through channel partners. It is updated twice a year and includes the latest Cisco offerings. To order and find out more about the Cisco Product Quick Reference Guide, go to this URL:

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

- 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

Send documentation comments to mdsfeedback-doc@cisco.com

Cisco Systems

Send documentation comments to mdsfeedback-doc@cisco.com

PART 1

Getting Started

CHAPTER 1

Overview

The Cisco MDS 9000 Family of multilayer directors and fabric switches offers intelligent fabric-switching services that realize maximum performance while ensuring high reliability levels. They combine robust and flexible hardware architecture with multiple layers of network and storage management intelligence. This powerful combination enables highly available, scalable storage networks that provide intelligent networking features such as multiprotocol and multitransport integration, virtual SANs (VSANs), advanced security, sophisticated debug analysis tools, and unified SAN management.

This chapter contains the following sections:

Hardware Overview

This section provides an overview of the following Cisco MDS 9000 Family of multilayer directors and fabric switches:

- Cisco MDS 9500 Series multilayer directors
  - Cisco MDS 9513 multilayer director
  - Cisco MDS 9509 multilayer director
  - Cisco MDS 9506 multilayer director
- Cisco MDS 9200 Series fabric switches
  - Cisco MDS 9216i multiprotocol fabric switch
  - Cisco MDS 9216A multilayer fabric switch
  - Cisco MDS 9216 multilayer fabric switch
- Cisco MDS 9100 Series fixed configuration fabric switches
  - Cisco MDS 9140 multilayer switch
  - Cisco MDS 9120 multilayer switch

Cisco MDS 9500 Series Multilayer Directors

The Cisco MDS 9500 Series includes the following multilayer, modular directors:

- The Cisco MDS 9513 Director, which has thirteen slots, two of which (slot 7 and slot 8) are reserved for the supervisor modules, and can accommodate up to eleven hot-pluggable switching or services modules.
- The Cisco MDS 9509 Director, which has nine slots, two of which (slot 5 and slot 6) are reserved for the supervisor modules, and can accommodate up to seven hot-pluggable switching or services modules.
- The Cisco MDS 9506 Director, which has six slots, two of which (slot 5 and slot 6) are reserved for the supervisor modules, and can accommodate up to four hot-pluggable switching or services modules.

Note Supervisor-1 modules and Supervisor-2 modules can only operate in the same chassis during migration.

The two supervisor modules ensure high availability and traffic load balancing capabilities. The standby supervisor module provides redundancy if the active supervisor module fails. The supervisor modules provide management access through a 10/100BASE-T Ethernet port switch and an RS-232 serial port.

Note The USB ports on the Supervisor-2 module are not supported by the Cisco MDS SAN-OS software.

The Cisco MDS 9500 Series directors support the following switching and services modules:

- 48-port 4-Gbps Fibre Channel switching module
- 24-port 4-Gbps Fibre Channel switching module
- 12-port 4-Gbps Fibre Channel switching module
- 4-port 10-Gbps Fibre Channel switching module
- 32-port 2-Gbps Fibre Channel switching module
- 16-port 2-Gbps Fibre Channel switching module
- 14/2-port Multiprotocol Services (MPS-14/2) module
- 8-port IP Storage Services (IPS-8) module
- 4-port IP Storage Services (IPS-4) module
- Storage Services Module (SSM)
- Caching Services Module (CSM)

Refer to the *Cisco MDS 9500 Series Hardware Installation Guide*.

Cisco MDS 9200 Series Fabric Switches

The Cisco MDS 9200 Series includes the following multilayer switches supporting multiprotocol capabilities:

- Cisco MDS 9216i
- Cisco MDS 9216A
- Cisco MDS 9216

Cisco MDS 9216i Multiprotocol Fabric Switch

The Cisco MDS 9216i multiprotocol fabric switch has two slots, one of which is reserved for the integrated supervisor module and the other for switching or services modules. The supervisor module provides supervisor functions and has 14 standard Fibre Channel ports and two multiprotocol ports that can support FCIP and iSCSI protocols simultaneously.

The Cisco MDS 9200 multilayer fabric switches support the following switching and services modules:

- 48-port 4-Gbps Fibre Channel switching module
- 24-port 4-Gbps Fibre Channel switching module
- 12-port 4-Gbps Fibre Channel switching module
- 4-port 10-Gbps Fibre Channel switching module
- 32-port 2-Gbps Fibre Channel switching module
- 16-port 2-Gbps Fibre Channel switching module
- 14/2-port Multiprotocol Services (MPS-14/2) module
- 8-port IP Storage Services (IPS-8) module
- 4-port IP Storage Services (IPS-4) module
- Storage Services Module (SSM)
- Caching Services Module (CSM)

Refer to the *Cisco MDS 9200 Series Hardware Installation Guide*.

Cisco MDS 9200 Multilayer Fabric Switches

The Cisco MDS 9216A and Cisco MDS 9216 multilayer fabric switches have two slots, one of which is reserved for the integrated supervisor module and the other for a switching or services module. The supervisor module provides supervisor functions and has 16 standard Fibre Channel ports.

The Cisco MDS 9216A multilayer fabric switch supports the following switching and services modules:

- 48-port 4-Gbps Fibre Channel switching module
- 24-port 4-Gbps Fibre Channel switching module
- 12-port 4-Gbps Fibre Channel switching module
- 4-port 10-Gbps Fibre Channel switching module
- 32-port 2-Gbps Fibre Channel switching module
- 16-port 2-Gbps Fibre Channel switching module
- 14/2-port Multiprotocol Services (MPS-14/2) module
- 8-port IP Storage Services (IPS-8) module
- 4-port IP Storage Services (IPS-4) module
- Caching Services Module (CSM)

The Cisco MDS 9216 multilayer fabric switch supports the following switching and services modules:

- 32-port 2-Gbps Fibre Channel switching module
- 16-port 2-Gbps Fibre Channel switching module
- 14/2-port Multiprotocol Services (MPS-14/2) module

- 8-port IP Storage Services (IPS-8) module
- 4-port IP Storage Services (IPS-4) module
- Caching Services Module (CSM)

Refer to the *Cisco MDS 9200 Series Hardware Installation Guide* and the *Cisco MDS 9216 Switch Hardware Installation Guide.*

Cisco MDS 9100 Series Fixed Configuration Fabric Switches

Cisco MDS 9100 Series includes two multilayer, fixed configuration (non-modular) switches:

- Cisco MDS 9140 with 40 ports (8 full-rate ports, 32 host-optimized ports)
- Cisco MDS 9120 with 20 ports (4 full-rate ports, 16 host-optimized ports)

These fixed configuration switches are packaged in 1 RU enclosures and provide 1-Gbps or 2-Gbps autosensing Fibre Channel ports. Besides Telnet access, a 10/100BASE-T Ethernet port provides switch access.

Note Switches in the Cisco MDS 9100 Series do not have a COM1 port (RS-232 serial port).

Refer to the *Cisco MDS 9100 Series Hardware Installation Guide*.

Cisco SAN-OS Software Configuration

This section describes the tools you can use to configure SAN-OS software, and provides an overview of the software configuration process with links to the appropriate chapters.

Note Fabric Manager also manages Cisco MDS 9020 switches running FabricWare 2.1. For more information, refer to the *Cisco MDS 9020 Fabric Switch Configuration Guide and Command Reference.*

This section includes the following topics:

Tools for Software Configuration

You can use one of two configuration management tools to configure your SANs (see Figure 1-1).

- The command-line interface (CLI) can manage Cisco MDS 9000 Family switches using Telnet, SSH, or a serial connection.
- The Cisco MDS 9000 Fabric Manager, a Java-based graphical user interface, can manage Cisco MDS 9000 Family switches using SNMP.

Figure 1-1 Tools for Configuring Cisco SAN-OS Software

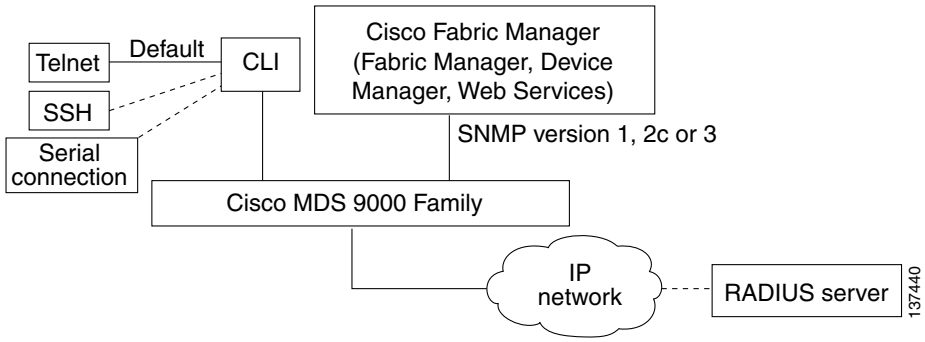

CLI

With the CLI, you can type commands at the switch prompt, and the commands are executed when you press the **Enter** key. The CLI parser provides command help, command completion, and keyboard sequences that allow you to access previously executed commands from the buffer history.

For more information on configuring the Cisco MDS switch using the CLI, refer to the *Cisco MDS 9000 CLI Configuration Guide*.

Cisco MDS 9000 Fabric Manager

The Cisco Fabric Manager is a set of network management tools that supports Secure Simple Network Management Protocol version 3 (SNMPv3) and legacy versions. The Cisco Fabric Manager applications are:

- Fabric Manager Client—provides a graphical user interface (GUI) that displays real-time views of your network fabric, and lets you manage the configuration of Cisco MDS 9000 Family devices and third-party switches.
- Fabric Manager Server—performs advanced monitoring, troubleshooting, and configuration for multiple fabrics. It must be started before running the Fabric Manager Client. It can be accessed by up to 16 Fabric Manager Clients at a time.
- Device Manager—presents two views of a switch.
  - Device View displays a continuously updated physical representation of the switch configuration, and provides access to statistics and configuration information for a single switch.
  - Summary View presents real-time performance statistics of all active interfaces and channels on the switch for Fibre Channel and IP connections.
- Fabric Manager Web Services—allows operators to monitor MDS events, performance, and inventory, and perform minor configuration tasks from a remote location using a web browser.
- Performance Manager—provides detailed traffic analysis by capturing data with SNMP. This data is compiled into various graphs and charts that can be viewed with any web browser using Fabric Manager Web Services.

The Cisco Fabric Manager applications are an alternative to the CLI for most switch configuration commands.

Note Resource Manager Essentials (RME) versions 3.4 and 3.5 provide support for switches in the Cisco MDS 9000 Family. Device Updates (DU) are available on Cisco.com (http://www.cisco.com/).

Continue reading this book for more information on configuring the Cisco MDS switch using the Cisco MDS 9000 Family Fabric Manager.

Software Configuration Overview

This section provides an overview of the SAN-OS configuration process and includes the following topics:

Basic Configuration

These sections contain the minimum information you need to get your switch up and running.

- Setting Up the Switch (Starting a Switch in the Cisco MDS 9000 Family, page 2-2)
- Installing Fabric Manager (Installing the Management Software, page 2-19)
  - Fabric Manager Server (Chapter 3, “Fabric Manager Server”)
  - Fabric Manager Client (Chapter 4, “Fabric Manager Client”)
  - Device Manager (Chapter 5, “Device Manager”)
  - Fabric Manager Web Services (Chapter 6, “Fabric Manager Web Services”)
- Installing Licenses (Chapter 10, “Obtaining and Installing Licenses”)
- Configuring the Minimum Requirements
  - Initial configuration (Chapter 11, “Initial Configuration”)
  - VSANs (Chapter 23, “Configuring and Managing VSANs”)
  - Interfaces (Chapter 18, “Configuring Interfaces”)
  - Zones and zone sets (Chapter 26, “Configuring and Managing Zones”)

Advanced Configuration

These sections contain additional configuration information for SAN-OS software and the MDS 9000 Family of switches, and includes the following topics:

Switch Configuration

- Generation 2 switching modules (Chapter 19, "Configuring Generation 2 Switching Modules")
- High Availability (Chapter 15, "Configuring High Availability")
- Trunking (Chapter 20, "Configuring Trunking")
- PortChannels (Chapter 21, "Configuring PortChannels")
- Domains Chapter 22, "Configuring Domain Parameters"

Fabric Configuration

- Dynamic VSANs (Chapter 24, "Creating Dynamic VSANs")
- Inter-VSAN Routing (Chapter 25, "Configuring Inter-VSAN Routing")
- Device alias distribution (Chapter 27, "Distributing Device Alias Services")
- FSPF Chapter 28, "Configuring Fibre Channel Routing Services and Protocols"
- FLOGI Chapter 29, "Managing FLOGI, Name Server, FDMI, and RSCN Databases"
- FICON (Chapter 31, "Configuring FICON")
- Switch interoperability (Chapter 32, "Advanced Features and Concepts")

Security

- Users and Roles (Chapter 33, "Configuring Users and Common Roles")
- SNMP (Chapter 34, "Configuring SNMP")
- RADIUS and TACACS+ (Chapter 35, "Configuring RADIUS and TACACS+")
- Access lists for IPv4 (Chapter 36, "Configuring IPv4 Access Control Lists")
- Access lists for IPv6 (Chapter 37, "Configuring IPv6 Access Control Lists")
- Digital certificates (Chapter 38, "Configuring Certificate Authorities and Digital Certificates")
- IPsec (Chapter 39, "Configuring IPsec Network Security")
- FC-SP (Chapter 40, "Configuring FC-SP and DHCHAP")
- Port security (Chapter 41, "Configuring Port Security")
- Fabric binding (Chapter 42, "Configuring Fabric Binding")

IP Services

- FCIP (Chapter 43, "Configuring FCIP")
- SAN extension tuner (Chapter 44, "Configuring the SAN Extension Tuner")
- iSCSI (Chapter 45, "Configuring iSCSI")
- IP Services (Chapter 46, "Configuring IP Services")
- IP Storage (Chapter 47, "Configuring IP Storage")
- IPv6 (Chapter 48, "Configuring IPv6 for Gigabit Ethernet Interfaces")

Intelligent Storage Services

- SCSI flow services (Chapter 49, "Configuring SCSI Flow Services and Statistics")

- Fibre Channel write acceleration (Chapter 50, “Configuring Fibre Channel Write Acceleration”
- SANTap (Chapter 51, “Configuring SANTap”)
- NASB (Chapter 52, “Configuring NASB”)

Network and Switch Monitoring

- General Network Monitoring (Chapter 53, “Network Monitoring”)
- Performance Monitoring (Chapter 54, “Performance Monitoring”)
- RMON (Chapter 55, “Configuring RMON”)
- SPAN (Chapter 56, “Monitoring Network Traffic Using SPAN”)
- System message logging (Chapter 57, “Configuring System Message Logging”)
- Call Home (Chapter 58, “Configuring Call Home”)
- Fabric configuration servers (Chapter 59, “Configuring Fabric Configuration Servers”)

Traffic Management

- QoS (Chapter 60, “Configuring Fabric Congestion Control and QoS”)
- Port Tracking (Chapter 61, “Configuring Port Tracking”)

CHAPTER 2

Installing Cisco MDS SAN-OS and Fabric Manager

Cisco Fabric Manager is a set of network management tools that supports Secure Simple Network Management Protocol version 3 (SNMPv3). It provides a graphical user interface (GUI) that displays real-time views of your network fabrics and lets you manage the configuration of Cisco MDS 9000 Family devices and third-party switches.

This chapter contains the following sections:

Starting a Switch in the Cisco MDS 9000 Family

The following procedure is a review of the tasks you should have completed during hardware installation, including starting up the switch. These tasks must be completed before you can configure the switch.

Note You must use the CLI for initial switch start up.

Before you can configure a switch, follow these steps:

Step 1 Verify the following physical connections for the new Cisco MDS 9000 Family switch:

- The console port is physically connected to a computer terminal (or terminal server).
- The management 10/100 Ethernet port (mgmt0) is connected to an external hub, switch, or router.

Refer to the *Cisco MDS 9000 Family Hardware Installation Guide* (for the required product) for more information.

Tip Save the host ID information for future use (for example, to enable licensed features). The host ID information is provided in the Proof of Purchase document that accompanies the switch.

Step 2 Verify that the default console port parameters are identical to those of the computer terminal (or terminal server) attached to the switch console port:

- 9600 baud
- 8 data bits
- 1 stop bit
- No parity

Step 3 Power on the switch. The switch boots automatically and the `switch#` prompt appears in your terminal window.

Initial Setup Routine

The first time that you access a switch in the Cisco MDS 9000 Family using the CLI, it runs a setup program that prompts you for the IP address and other configuration information necessary for the switch to communicate over the supervisor module Ethernet interface. This information is required to configure and manage the switch.

Note The IP address can only be configured from the CLI. When you power up the switch for the first time, assign the IP address. After you perform this step, the Cisco MDS 9000 Family Fabric Manager can reach the switch through the console port.

Preparing to Configure the Switch

Before you configure a switch in the Cisco MDS 9000 Family for the first time, you need the following information:

- Administrator password, including:
  - Creating a password for the administrator (required).
  - Creating an additional login account and password (optional).
- IP address for the switch management interface—The management interface can be an out-of-band Ethernet interface or an in-band Fibre Channel interface (recommended).
- Subnet mask for the switch's management interface (optional).
- IP addresses, including:
  - Destination prefix, destination prefix subnet mask, and next hop IP address, if you want to enable IP routing. Also, provide the IP address of the default network (optional).
  - You need an IP address of the default gateway (optional).
- SSH service on the switch—To enable this optional service, select the type of SSH key (dsa/rsa/rsa1) and number of key bits (768 to 2048).
- DNS IP address (optional).
- Default domain name (optional).
- NTP server IP address (optional).
- SNMP community string (optional).
- Switch name—This name is your switch prompt (optional).

Note Be sure to configure the IP route, the IP default network address, and the IP default gateway address to enable SNMP access. If IP routing is enabled, the switch uses the IP route and the default network IP address. If IP routing is disabled, the switch uses the default gateway IP address.

Default Login

All Cisco MDS 9000 Family switches have the network administrator as a default user (admin). You cannot change the default user at any time (see the "Role-Based Authorization" section on page 33-1).

You must explicitly configure a strong password for any switch in the Cisco MDS 9000 Family. If a password is trivial (short, easy to decipher), your password configuration is rejected. Be sure to configure a strong password (see the "User Accounts" section on page 33-10). If you configure and then forget this new password, you can recover this password (see the "Recovering the Administrator Password" section on page 33-19).

Setup Options

The setup scenario differs based on the subnet to which you are adding the new switch. You must configure a Cisco MDS 9000 Family switch with an IP address to enable management connections from outside of the switch. The setup options are as follows:

- Out-of-band management—This feature provides a connection to the network through a supervisor module front panel Ethernet port (see Figure 2-1).
- In-band management—This feature provides IP over Fibre Channel (IPFC) to manage the switches. The in-band management feature is transparent to the network management system (NMS). Instead of conventional Ethernet physical media, switches in the Cisco MDS 9000 Family use IPFC as the transport mechanism (see Figure 2-1 and Chapter 46, "Configuring IP Services").

Figure 2-1 Management Access to Switches

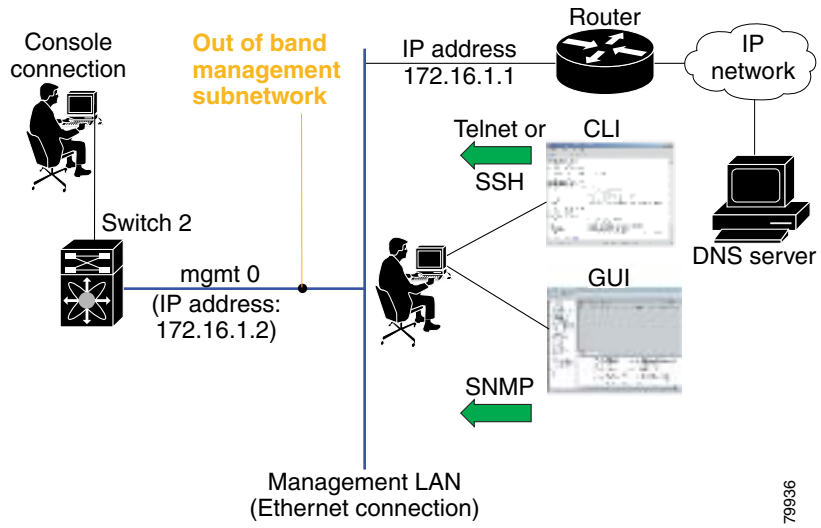

Assigning Setup Information

This section describes how to initially configure the switch for both out-of-band and in-band management.

Note Press **Ctrl-C** at any prompt to skip the remaining configuration options and proceed with what is configured until that point. Entering the new password for the administrator is a requirement and cannot be skipped.

Tip If you do not wish to answer a previously configured question, or if you wish to skip answers to any questions, press **Enter**. If a default answer is not available (for example, the switch name), the switch uses what was previously configured and skips to the next question.

Configuring Out-of-Band Management

Note You can configure both in-band and out-of-band configuration together by entering **Yes** in both Step 11c and Step 11d in the following procedure.

To configure the switch for first time out-of-band access, follow these steps:

Step 1 Power on the switch. The switches in the Cisco MDS 9000 Family boot automatically.

Step 2 Enter the new password for the administrator.

```
Enter the password for admin: 2004asdf*lkjh18
```

Tip If a password is trivial (short, easy to decipher), your password configuration is rejected. Be sure to configure a strong password as shown in the sample configuration. Passwords are case-sensitive. You must explicitly configure a password that meets the requirements listed in the "User Accounts" section on page 33-10.

Step 3 Enter **yes** to enter the setup mode.

```
This setup utility will guide you through the basic configuration of the system. Setup
configures only enough connectivity for management of the system.

Please register Cisco MDS 9000 Family devices promptly with your supplier. Failure to
register may affect response times for initial service calls. MDS devices must be
registered to receive entitled support services.

Press Enter incase you want to skip any dialog. Use ctrl-c at anytime to skip away
remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): yes
```

The setup utility guides you through the basic configuration process. Press **Ctrl-C** at any prompt to end the configuration process.

Step 4 Enter the new password for the administrator (admin is the default).

```
Enter the password for admin: admin
```

Step 5 Enter **yes** (no is the default) to create additional accounts.

```
Create another login account (yes/no) [n]: yes
```

While configuring your initial setup, you can create an additional user account (in the network-admin role) in addition to the administrator's account. See the "Role-Based Authorization" section on page 33-1 for information on default roles and permissions.

Note User login IDs must contain non-numeric characters.

a. Enter the user login ID.

```
Enter the user login ID: user_name
```

b. Enter the user password.

```
Enter the password for user_name: user-password
```

Step 6 Enter **yes** (yes is the default) to create an SNMPv3 account.

```
Configure SNMPv3 Management parameters (yes/no) [y]: yes
```

a. Enter the user name (admin is the default).

```
SNMPv3 user name [admin]: admin
```

b. Enter the SNMPv3 password (minimum of eight characters). The default is **admin123**.

```
SNMPv3 user authentication password: admin_pass
```

Note By default, if the admin password is at least eight characters, then the SNMP authentication password is the same as the admin password (at least eight characters). If the admin password is less than eight characters, then you need to provide a new password for SNMP.
The admin password can have a minimum of one character, but the SNMP authentication password must have a minimum of eight characters.

Step 7 Enter **yes** (no is the default) to configure the read-only or read-write SNMP community string.

```
Configure read-only SNMP community string (yes/no) [n]: yes
```

a. Enter the SNMP community string.

```
SNMP community string: snmp_community
```

Step 8 Enter a name for the switch.

Note The switch name is limited to 32 alphanumeric characters. The default is **switch**.

```
Enter the switch name: switch_name
```

Step 9 Enter **yes** (yes is the default) to configure out-of-band management.

```
Continue with Out-of-band (mgmt0) management configuration? [yes/no]: yes
```

a. Enter the mgmt0 IP address.

```
Mgmt0 IP address: ip_address
```

b. Enter the mgmt0 subnet mask.

```
Mgmt0 IP netmask: subnet_mask
```

Step 10 Enter **yes** (yes is the default) to configure the default gateway (recommended).

```
Configure the default-gateway: (yes/no) [y]: yes
```

a. Enter the default gateway IP address.

```
IP address of the default-gateway: default_gateway
```

Step 11 Enter **yes** (**no** is the default) to configure advanced IP options such as in-band management, static routes, default network, DNS, and domain name.

```
Configure Advanced IP options (yes/no)? [n]: yes
```

a. Enter **no** (no is the default) at the in-band management configuration prompt.

```
Continue with in-band (VSAN1) management configuration? (yes/no) [no]: no
```

b. Enter **yes** (yes is the default) to enable IP routing capabilities.

```
Enable the ip routing? (yes/no) [y]: yes
```

c. Enter **yes** (yes is the default) to configure a static route (recommended).

```
Configure static route: (yes/no) [y]: yes
```

Enter the destination prefix.

```
Destination prefix: dest_prefix
```

Type the destination prefix mask.

```
Destination prefix mask: dest_mask
```

Type the next hop IP address.

```
Next hop ip address: next_hop_address
```

Note Be sure to configure the IP route, the default network IP address, and the default gateway IP address to enable SNMP access. If IP routing is enabled, the switch uses the IP route and the default network IP address. If IP routing is disabled, the switch uses the default gateway IP address.

d. Enter **yes** (yes is the default) to configure the default network (recommended).

```
Configure the default network: (yes/no) [y]: yes
```

Enter the default network IP address.

Note The default network IP address is the destination prefix provided in Step 11c .

```
Default network IP address [dest_prefix]: dest_prefix
```

e. Enter **yes** (yes is the default) to configure the DNS IP address.

```
Configure the DNS IP address? (yes/no) [y]: yes
```

Enter the DNS IP address.

```
DNS IP address: name_server
```

f. Enter **yes** (default is no) to configure the default domain name.

```
Configure the default domain name? (yes/no) [n]: yes
```

Enter the default domain name.

```
Default domain name: domain_name
```

Step 12 Enter **yes** (yes is the default) to enable Telnet service.

```
Enable the telnet service? (yes/no) [y]: yes
```

Step 13 Enter **yes** (no is the default) to enable the SSH service.

```
Enabled SSH service? (yes/no) [n]: yes
```

Step 14 Enter the SSH key type (see the "Generating the SSH Server Key Pair" section on page 33-16) that you would like to generate.

```
Type the SSH key you would like to generate (dsa/rsa/rsa1)? dsa
```

Step 15 Enter the number of key bits within the specified range.

```
Enter the number of key bits? (768 to 2048): 768
```

Step 16 Enter **yes** (no is the default) to configure the NTP server.

```
Configure NTP server? (yes/no) [n]: yes
```

a. Enter the NTP server IP address.

```
NTP server IP address: ntp_server_IP_address
```

Step 17 Enter **noshut** (shut is the default) to configure the default switch port interface to the shut state.

```
Configure default switchport interface state (shut/noshut) [shut]: noshut
```

Step 18 Enter **on** (on is the default) to configure the switch port trunk mode.

```
Configure default switchport trunk mode (on/off/auto) [on]: on
```

Step 19 Enter **on** (off is the default) to configure the PortChannel auto-create state.

```
Configure default port-channel auto-create state (on/off) [off]: on
```

Step 20 Enter **permit** (deny is the default) to deny a default zone policy configuration.

```
Configure default zone policy (permit/deny) [deny]: permit
```

Permits traffic flow to all members of the default zone.

Step 21 Enter **yes** (no is the default) to disable a full zone set distribution (see the "Zone Set Distribution" section on page 26-24).

```
Enable full zoneset distribution (yes/no) [n]: yes
```

Disables the switch-wide default for the full zone set distribution feature.

You see the new configuration. Review and edit the configuration that you have just entered.

Step 22 Enter **no** (no is the default) if you are satisfied with the configuration.

```
The following configuration will be applied:
  username admin password admin_pass role network-admin
  username user_name password user_pass role network-admin
  snmp-server community snmp_community ro
  switchname switch
  interface mgmt0
    ip address ip_address subnet_mask
    no shutdown
  ip routing
  ip route dest_prefix dest_mask dest_address
  ip default-network dest_prefix
  ip default-gateway default_gateway
  ip name-server name_server
  ip domain-name domain_name
  telnet server enable
  ssh key dsa 768 force
  ssh server enable
  ntp server ipaddr ntp_server
  system default switchport shutdown
  system default switchport trunk mode on
  system default port-channel auto-create
  zone default-zone permit vsan 1-4093
  zoneset distribute full vsan 1-4093

Would you like to edit the configuration? (yes/no) [n]: no
```

Step 23 Enter **yes** (yes is the default) to use and save this configuration:

```
Use this configuration and save it? (yes/no) [y]: yes
```

Caution If you do not save the configuration at this point, none of your changes are updated the next time the switch is rebooted. Type **yes** to save the new configuration. This action ensures that the kickstart and system images are also automatically configured (see Chapter 13, "Software Images").

Configuring In-Band Management

The in-band management logical interface is VSAN 1. This management interface uses the Fibre Channel infrastructure to transport IP traffic. An interface for VSAN 1 is created on every switch in the fabric. Each switch should have its VSAN 1 interface configured with an IP address in the same subnetwork. A default route that points to the switch that provides access to the IP network should be configured on every switch in the Fibre Channel fabric (see Chapter 23, "Configuring and Managing VSANs").

Note You can configure both in-band and out-of-band configuration together by entering **Yes** in both Step 9c and Step 9d in the following procedure.

To configure a switch for first time in-band access, follow these steps:

Step 1 Power on the switch. Switches in the Cisco MDS 9000 Family boot automatically.

Step 2 Enter the new password for the administrator.

```
Enter the password for admin: 2004asdf*lkjh18
```

Tip If a password is trivial (short, easy to decipher), your password configuration is rejected. Be sure to configure a strong password as shown in the sample configuration. Passwords are case-sensitive. You must explicitly configure a password that meets the requirements listed in the "User Accounts" section on page 33-10.

Step 3 Enter **yes** to enter the setup mode.

```
This setup utility will guide you through the basic configuration of the system. Setup
configures only enough connectivity for management of the system.

Please register Cisco MDS 9000 Family devices promptly with your supplier. Failure to
register may affect response times for initial service calls. MDS devices must be
registered to receive entitled support services.

Press Enter incase you want to skip any dialog. Use ctrl-c at anytime to skip away
remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): yes
```

The setup utility guides you through the basic configuration process. Press **Ctrl-C** at any prompt to end the configuration process.

Step 4 Enter **no** (no is the default) if you do not wish to create additional accounts.

```
Create another login account (yes/no) [no]: no
```

Step 5 Configure the read-only or read-write SNMP community string.

a. Enter **no** (no is the default) to avoid configuring the read-only SNMP community string.

```
Configure read-only SNMP community string (yes/no) [n]: no
```

b. Enter **no** (no is the default) to configure the read-only SNMP community string.

```
Configure read-only SNMP community string (yes/no) [n]: yes
```

c. Enter the SNMP community string.

```
SNMP community string: snmp_community
```

Step 6 Enter a name for the switch.

Note The switch name is limited to 32 alphanumeric characters. The default is **switch**.

```
Enter the switch name: switch_name
```

Step 7 Enter **no** (yes is the default) at the configuration prompt to configure out-of-band management.

```
Continue with Out-of-band (mgmt0) management configuration? [yes/no]: no
```

Step 8 Enter **yes** (yes is the default) to configure the default gateway.

```
Configure the default-gateway: (yes/no) [y]: yes
```

a. Enter the default gateway IP address.

```
IP address of the default gateway: default_gateway
```

Step 9 Enter **yes** (**no** is the default) to configure advanced IP options such as in-band management, static routes, default network, DNS, and domain name.

```
Configure Advanced IP options (yes/no)? [n]: yes
```

a. Enter **yes** (no is the default) at the in-band management configuration prompt.

```
Continue with in-band (VSAN1) management configuration? (yes/no) [no]: yes
```

Enter the VSAN 1 IP address.

```
VSAN1 IP address: ip_address
```

Enter the subnet mask.

```
VSAN1 IP net mask: subnet_mask
```

b. Enter **no** (yes is the default) to enable IP routing capabilities.

```
Enable ip routing capabilities? (yes/no) [y]: no
```

c. Enter **no** (yes is the default) to configure a static route.

```
Configure static route: (yes/no) [y]: no
```

d. Enter **no** (yes is the default) to configure the default network.

```
Configure the default-network: (yes/no) [y]: no
```

e. Enter **no** (yes is the default) to configure the DNS IP address.

```
Configure the DNS IP address? (yes/no) [y]: no
```

f. Enter **no** (no is the default) to skip the default domain name configuration.

```
Configure the default domain name? (yes/no) [n]: no
```

Step 10 Enter **no** (yes is the default) to disable Telnet service.

```
Enable the telnet service? (yes/no) [y]: no
```

Step 11 Enter **yes** (no is the default) to enable the SSH service.

```
Enabled SSH service? (yes/no) [n]: yes
```

Step 12 Enter the SSH key type (see the "Generating the SSH Server Key Pair" section on page 33-16) that you would like to generate.

```
Type the SSH key you would like to generate (dsa/rsa/rsa1)? rsa
```

Step 13 Enter the number of key bits within the specified range.

```
Enter the number of key bits? (768 to 1024): 1024
```

Step 14 Enter **no** (no is the default) to configure the NTP server.

```
Configure NTP server? (yes/no) [n]: no
```

Step 15 Enter **shut** (shut is the default) to configure the default switch port interface to the shut state.

```
Configure default switchport interface state (shut/noshut) [shut]: shut
```

Note The management Ethernet interface is not shut down at this point—only the Fibre Channel, iSCSI, FCIP, and Gigabit Ethernet interfaces are shut down.

Step 16 Enter **auto** (off is the default) to configure the switch port trunk mode.

```
Configure default switchport trunk mode (on/off/auto) [off]: auto
```

Step 17 Enter **off** (off is the default) to configure the PortChannel auto-create state.

```
Configure default port-channel auto-create state (on/off) [off]: off
```

Step 18 Enter **deny** (deny is the default) to deny a default zone policy configuration.

```
Configure default zone policy (permit/deny) [deny]: deny
```

Denies traffic flow to all members of the default zone.

Step 19 Enter **no** (no is the default) to disable a full zone set distribution (see the "Zone Set Distribution" section on page 26-24).

```
Enable full zoneset distribution (yes/no) [n]: no
```

Disables the switch-wide default for the full zone set distribution feature.

You see the new configuration. Review and edit the configuration that you have just entered.

Step 20 Enter **no** (no is the default) if you are satisfied with the configuration.

```
The following configuration will be applied:
  username admin password admin_pass role network-admin
  snmp-server community snmp_community rw
  switchname switch
  interface vsan1
    ip address ip_address subnet_mask
    no shutdown
  ip default-gateway default_gateway
  no telnet server enable
  ssh key rsa 1024 force
  ssh server enable
  no system default switchport shutdown
  system default switchport trunk mode auto
  no zone default-zone permit vsan 1-4093
  no zoneset distribute full vsan 1-4093

Would you like to edit the configuration? (yes/no) [n]: no
```

Step 21 Enter **yes** (yes is default) to use and save this configuration.

```
Use this configuration and save it? (yes/no) [y]: yes
```

Caution If you do not save the configuration at this point, none of your changes are updated the next time the switch is rebooted. Type **yes** to save the new configuration. This ensures that the kickstart and system images are also automatically configured (see Chapter 13, "Software Images").

Using the setup Command

To make changes to the initial configuration at a later time, you can issue the **setup** command in EXEC mode.

```
switch# setup
---- Basic System Configuration Dialog ----
This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.
*Note: setup always assumes a predefined defaults irrespective
of the current system configuration when invoked from CLI.

Press Enter incase you want to skip any dialog. Use ctrl-c at anytime
to skip away remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): yes
```

The setup utility guides you through the basic configuration process.

Accessing the Switch

After initial configuration, you can access the switch in one of three ways (see Figure 2-2):

- Serial console access—You can use a serial port connection to access the CLI.
- In-band IP (IPFC) access—You can use Telnet or SSH to access a switch in the Cisco MDS 9000 Family or use Cisco MDS 9000 Fabric Manager to access the switch.
- Out-of-band (10/100BASE-T Ethernet) access—You can use Telnet or SSH to access a switch in the Cisco MDS 9000 Family or use Cisco MDS 9000 Fabric Manager to access the switch.

Figure 2-2 Switch Access Options

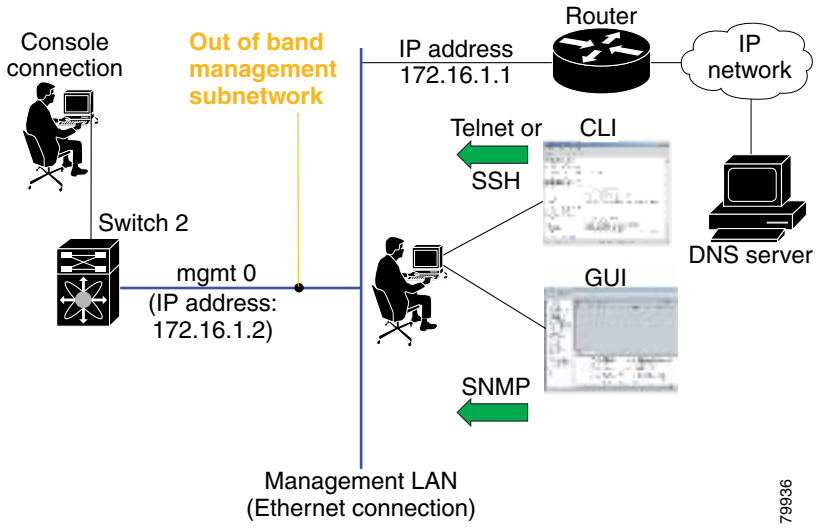

Where Do You Go Next?

After reviewing the default configuration, you can change it or perform other configuration or management tasks. The initial setup can only be performed at the CLI. However, you can continue to configure other software features, or access the switch after initial configuration by using either the CLI or the Device Manager and Fabric Manager applications.

To use the CLI, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

About Cisco Fabric Manager

The Cisco Fabric Manager provides an alternative to the command-line interface (CLI) for most switch configuration commands. For information on using the CLI to configure a Cisco MDS 9000 Family switch, refer to the *Cisco MDS 9000 Family CLI Configuration Guide* or the *Cisco MDS 9020 Switch Configuration Guide and Command Reference Guide.* For details on managing switches running Cisco FabricWare, see the "Managing Cisco FabricWare with Fabric Manager" section on page C-3.

In addition to complete configuration and status monitoring capabilities for Cisco MDS 9000 switches, Fabric Manager provides powerful Fibre Channel troubleshooting tools. These in-depth health and configuration analysis capabilities leverage unique MDS 9000 switch capabilities: Fibre Channel Ping and Traceroute.

The Cisco Fabric Manager includes these management applications:

- Fabric Manager (client and server)
- Device Manager
- Performance Manager
- Fabric Manager Web Services

Fabric Manager Server

The Fabric Manager Server component must be started before running Fabric Manager. On a Windows PC, the Fabric Manager Server is installed as a service. This service can then be administered using the Windows Services in the Control Panel. Fabric Manager Server is responsible for discovery of the physical and logical fabric, and for listening for SNMP traps, syslog messages, and Performance Manager threshold events. For more information, see Chapter 3, "Fabric Manager Server."

Fabric Manager Client

The Fabric Manager Client component displays a map of your network fabrics, including Cisco MDS 9000 Family switches, third-party switches, hosts, and storage devices. The Fabric Manager Client provides multiple menus for accessing the features of the Fabric Manager Server. For more information, see Chapter 4, "Fabric Manager Client."

Fabric Manager Server Proxy Services

The Fabric Manager Client and Device Manager use SNMP to communicate with the Fabric Manager Server. In typical configurations, the Fabric Manager Server may be installed behind a firewall. The SNMP proxy service available in Cisco Fabric Manager Release 2.1(1a) or later provides a TCP-based transport proxy for these SNMP requests. The SNMP proxy service allows you to block all UDP traffic at the firewall and configure Fabric Manager Client to communicate over a configured TCP port.

Fabric Manager uses the CLI for managing some features on the switches. These management tasks are used by Fabric Manager and do not use the proxy services. Your firewall must remain open for CLI access for the following:

- External and internal loopback test
- Flash files
- Create CLI user
- Security - ISCSI users
- Quiese PC
- Show image version
- Show tech
- Switch resident reports (syslog, accounting)
- Zone migration
- Show cores

If you are using the SNMP proxy service and another application on your server is using port 8080, you need to modify your workstation settings.

To modify a Windows workstation, follow these steps:

Step 1 Open Internet Explorer and select **Tools > Internet Options**.

You see the Internet Options dialog box.

Step 2 Select the **Connections** tab and click **LAN Settings**.

You see the LAN Settings dialog box.

Step 3 Check the **Use a Proxy Server for your LAN** check box and click **Advanced**.

Step 4 Add your server IP Address or local host under the Exceptions section.

Step 5 Click **OK** to save your changes.

See the "Running Fabric Manager Behind a Firewall" section on page 2-25.

Device Manager

The Device Manager presents two views of a single switch.

- Device View displays a graphic representation of the switch configuration and provides access to statistics and configuration information.
- Summary View displays a summary of xE ports (Inter-Switch Links), Fx ports (fabric ports), and Nx ports (attached hosts and storage) on the switch, as well as Fibre Channel and IP neighbor devices. Summary or detailed statistics can be charted, printed, or saved to a file in tab-delimited format.

See Chapter 5, "Device Manager."

Performance Manager

Performance Manager provides detailed traffic analysis by capturing data with SNMP. This data is compiled into various graphs and charts that can be viewed with any web browser. See Chapter 54, "Performance Monitoring."

Fabric Manager Web Services

The Fabric Manager Web Services allows operators to monitor and obtain reports for MDS events, performance, and inventory from a remote location using a web browser. For information on installing and using Fabric Manager Web Services, see Chapter 6, "Fabric Manager Web Services."

Cisco MDS 9000 Switch Management

The Cisco MDS 9000 Family of switches can be accessed and configured in many different ways and supports standard management protocols. Table 2-1 lists the management protocols that Fabric Manager supports to access, monitor, and configure the Cisco MDS 9000 Family of switches .

Table 2-1 Supported Management Protocols

Management Protocol	Purpose
Telnet/SSH	Provides remote access to the CLI for a Cisco MDS 9000 switch.
FTP/SFTP/TFTP, SCP	Copies configuration and software images between devices.
SNMPv1, v2c, and v3	Includes over 80 distinct Management Information Bases (MIBs). Cisco MDS 9000 Family switches support SNMP version 1, 2, and 3 and RMON V1 and V2. RMON provides advanced alarm and event management, including setting thresholds and sending notifications based on changes in device or network behavior. By default, the Cisco Fabric Manager communicates with Cisco MDS 9000 Family switches using SNMPv3, which provides secure authentication using encrypted user names and passwords. SNMPv3 also provides the option to encrypt all management traffic.
HTTP/HTTPS	Includes HTTP and HTTPS for web browsers to communicate with Fabric Manager Web Services and for the distribution and installation of the Cisco Fabric Manager software. It is not used for communication between the Cisco Fabric Manager Server and Cisco MDS 9000 Family switches.
XML/CIM over HTTP/HTTPS	Includes CIM server support for designing storage area network management applications to run on Cisco SAN-OS.
ANSI T11 FC-GS-3	Provides Fibre Channel-Generic Services (FC-GS-3) in the defining management servers in the Fabric Configuration Server (FCS). Fabric Manager uses the information provided by FCS on top of the information contained in the Name Server database and in the Fibre Channel Shortest Path First (FSPF) topology database to build a detailed topology view and collect information for all the devices building the fabric.

Storage Management Solutions Architecture

Management services required for the storage environment can be divided into five layers, with the bottom layer being closest to the physical storage network equipment, and the top layer managing the interface between applications and storage resources.

Of these five layers of storage network management, Cisco Fabric Manager provides tools for device (element) management and fabric management. In general, the Device Manager is most useful for device management (a single switch), while Fabric Manager is more efficient for performing fabric management operations involving multiple switches.

Tools for upper-layer management tasks can be provided by Cisco or by third-party storage and network management applications. The following summarizes the goals and function of each layer of storage network management:

- Device management provides tools to configure and manage a device within a system or a fabric. You use device management tools to perform tasks on one device at a time, such as initial device configuration, setting and monitoring thresholds, and managing device system images or firmware.
- Fabric management provides a view of an entire fabric and its devices. Fabric management applications provide fabric discovery, fabric monitoring, reporting, and fabric configuration.
- Resource management provides tools for managing resources such as fabric bandwidth, connected paths, disks, I/O operations per second (IOPS), CPU, and memory. You can use Fabric Manager to perform some of these tasks.
- Data management provides tools for ensuring the integrity, availability, and performance of data. Data management services include redundant array of independent disks (RAID) schemes, data replication practices, backup or recovery requirements, and data migration. Data management capabilities are provided by third-party tools.
- Application management provides tools for managing the overall system consisting of devices, fabric, resources, and data from the application. Application management integrates all these components with the applications that use the storage network. Application management capabilities are provided by third-party tools.

In-Band Management and Out-of-Band Management

Cisco Fabric Manager requires an out-of-band (Ethernet) connection to at least one Cisco MDS 9000 Family switch. You need either mgmt0 or IP over Fibre Channel (IPFC) to manage the fabric.

mgmt0

The out-of-band management connection is a 10/100 Mbps Ethernet interface on the supervisor module, labeled mgmt0. The mgmt0 interface can be connected to a management network to access the switch through IP over Ethernet. You must connect to at least one Cisco MDS 9000 Family switch in the fabric through its Ethernet management port. You can then use this connection to manage the other switches using in-band (Fibre Channel) connectivity. Otherwise, you need to connect the mgmt0 port on each switch to your Ethernet network.

Each supervisor module has its own Ethernet connection; however, the two Ethernet connections in a redundant supervisor system operate in active or standby mode. The active supervisor module also hosts the active mgmt0 connection. When a failover event occurs to the standby supervisor module, the IP address and media access control (MAC) address of the active Ethernet connection are moved to the standby Ethernet connection.

IPFC

You can also manage switches on a Fibre Channel network using an in-band IP connection. The Cisco MDS 9000 Family supports RFC 2625 IP over Fibre Channel, which defines an encapsulation method to transport IP over a Fibre Channel network.

IPFC encapsulates IP packets into Fibre Channel frames so that management information can cross the Fibre Channel network without requiring a dedicated Ethernet connection to each switch. This feature allows you to build a completely in-band management solution.

Installing the Management Software

To install the software for the first time, or if you want to update or reinstall the software, access the supervisor module with a web browser. Click the **Install** links on the web page that is displayed. The software running on your workstation is verified to make sure you are running the most current version of the software. If it is not current, the most recent version is downloaded and installed on your workstation.

Note Before upgrading or uninstalling Fabric Manager or Device Manager, make sure any instances of these applications have been shut down.

Installation options include:

- Upgrade/Downgrade - The installer detects your current version of Fabric Manager and Device Manager, and it provides the option to upgrade or downgrade. The default is to upgrade to the latest version of Fabric Manager or Device Manager.
- Uninstall - If you are downgrading from Fabric Manager 2.x or later to Fabric Manager 1.3x or earlier, use the Uninstall batch file or shell script. Do not delete the MDS 9000 folder as this might prevent your installation from being upgraded in the future.

Note We recommend that you install the latest version of the Fabric Manager applications. Fabric Manager is backward-compatible with the Cisco MDS SAN-OS and Cisco FabricWare software running on the switches. When upgrading, upgrade the Fabric Manager software first, and then upgrade the Cisco MDS SAN-OS or Cisco FabricWare software on the switch.

Before You Install

Before you can access the Cisco Fabric Manager, you must complete the following tasks:

- Install a supervisor module on each switch that you want to manage.
- Configure the supervisor module with the following values using the setup routine or the CLI:
  - IP address assigned to the mgmt0 interface
  - SNMP credentials (v3 user name and password or v1/v2 communities), maintaining the same user name and password for all the switches in the fabric

Cisco MDS SAN-OS Release 2.1(1a) or later supports AAA authentication using RADIUS, TACACS+, or local SNMP users.

The Cisco Fabric Manager software executable files reside on each supervisor module of each Cisco MDS 9000 Family switch running Cisco MDS SAN-OS software in your network. The supervisor module provides an HTTP server that responds to browser requests and distributes the software to Windows or UNIX network management stations. You can also find Cisco Fabric Manager software on Cisco.com at the following website:

http://cisco.com/cgi-bin/tablebuild.pl/mds-fm

The Cisco Fabric Manager management software has been tested with the following software:

- Operating Systems
  - Windows 2000, 2003, XP
  - Solaris 2.8
  - Redhat Linux 7.2
- Java
  - Sun JRE and JDK 1.4.0, 1.4.1, 1.4.2, and 1.5.0 (recommended)
  - Java Web Start 1.2 and 1.0.1
- Browsers
  - Internet Explorer 5.5 or later
  - Netscape 6 or later
  - Mozilla 1.0 or later

Installation Procedure

For switches running Cisco MDS 9000 FabricWare, you need to install the Fabric Manager software from the CD-ROM included with your switch, or download Fabric Manager from Cisco.com.

To install Fabric Manager from the CD-ROM, navigate to the Fabric Manager installation notes and follow the directions.

To download the software from Cisco.com, go to the following website:

http://cisco.com/cgi-bin/tablebuild.pl/mds-fm

To install Fabric Manager on your workstation, follow these steps:

Step 1 Optionally, enter the IP address or host name of the supervisor module running Cisco MDS SAN-OS in the Address or Location field of your browser.

When you connect to the server for the first time, it checks to see if you have the correct Sun Java Virtual Machine version installed on your workstation. If not, a link is provided to the appropriate Sun Microsystems web page so you can install it. Fabric Manager looks for version 1.4(x) during installation.

The supervisor module HTTP server displays the installation window.

Step 2 Click the link to the Sun Java Virtual Machine software (if required) and install the software.

Using the instructions provided by the Sun Microsystems website, reconnect to the supervisor module by reentering the IP address or host name in the Location or Address field of your browser.

Note In some cases, license validation from Cisco partners requires Java version 1.4.2_04 or later. If you cannot install licenses from a Cisco partner, check to make sure your Java version is at least 1.4.2_04.

Note You can run CiscoWorks on the same PC as Fabric Manager, even though the Java requirements are different. When installing the later Java version for Fabric Manager, make sure it does not overwrite the earlier Java version required for CiscoWorks. Both versions of Java can coexist on your PC.

Step 3 Click the desired installation link (**Fabric Manager**, **Device Manager**, or **Fabric Manager Web Services and Performance Manager**).

Step 4 Select an installation folder for Fabric Manager on your workstation, as shown in Figure 2-3. The default location is C:\Program Files\Cisco Systems\MDS 9000 for Windows. On a UNIX (Solaris or Linux) machine, the installation path name is /usr/local/cisco_mds9000 or $HOME/cisco_mds9000, depending on the permissions of the user doing the installation.

Step 5 Check the Use Global Aliases in place of FC Aliases check box if you want to use global device aliases or replace existing per VSAN FC aliases with global device aliases.

Figure 2-3 Install Options

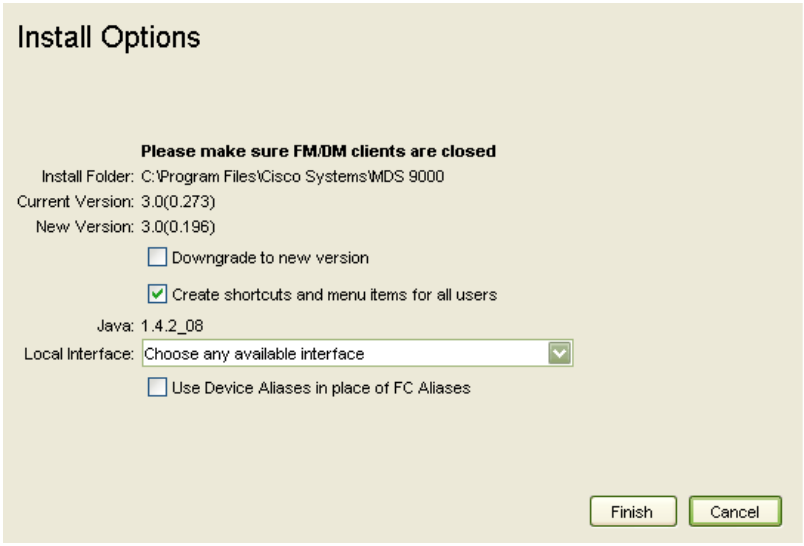

Tip After installation, you can choose to use global aliases by setting fabric.globalAlias to **true** in the server.properties file. In Fabric Manager Release 2.1(2) or later, you can select **Server > Admin** and check the **Device Alias** check box to use global aliases, or you can uncheck **Device Alias** to use FC aliases.

Step 6 Check the **Don't install and run FM Server** check box if you are installing just the Fabric Manager Client on a remote workstation.

A Cisco MDS 9000 program group is created under Start > Programs on Windows. This program group contains shortcuts to batch files in the install directory. Three services are started: Fabric Manager Server, Database, and Web Server. The Performance Manager server is installed but the service is not started at install time, as certain setup steps must be completed first.

On a UNIX (Solaris or Linux) machine, shell scripts are created in the install directory. The shell scripts that run the programs equivalent to the Windows services are: FMServer.sh, FMPersist.sh, PMCollector.sh, and FMWebClient.sh. All server-side data and Performance Manager data are stored under the install directory.

Fabric Manager Client cannot run without the server component, Fabric Manager Server. The server component is downloaded and installed when you download and install Fabric Manager or Device Manager. On a Windows machine you install the Fabric Manager Server as a service. This service can then be administered using Services in the Microsoft Windows Control Panel. The default setting for the Fabric Manager Server service is that the server is automatically started when the machine is rebooted. You can change this behavior by modifying the properties in Services.

Upgrading the Management Software

If you log into a switch running Cisco MDS SAN-OS with Fabric Manager or Device Manager, and that switch has a later version of the management software, you are prompted to install the later version. To upgrade the Cisco MDS Fabric Manager software, follow the instructions described in the "Installing the Management Software" section on page 2-19. You can also upgrade the software at any time by entering the IP address or host name of the supervisor module with the later version of software in the Address (Location) field of your browser.

Downgrading the Management Software

Before downgrading Fabric Manager, check to see which versions of Cisco SAN-OS are runnning on the switches in the fabric. If the switches are running a version of SAN-OS later than the downgraded Fabric Manager version, you will not be able to configure the newer features on those switches. For example, if the switches are running SAN-OS 3.x and you downgrade Fabric Manager to version 2.1, you will not be able to configure any of the 3.x features on those switches.

Downgrading to Release 2.x Versions

You can downgrade to Release 2.x from a newer release of Fabric Manager without removing the old version. For example, you can downgrade from:

- Fabric Manager Release 3.x to any version of 2.x
- Fabric Manager Release 2.x to an earlier version of 2.x

To downgrade, follow these steps:

Step 1 Close all instances of Fabric Manager Client or Device Manager on your workstation.

Step 2 Enter the IP address or host name of the supervisor module with the lower version of software in the Address or Location field of your browser and follow the installation steps. For more information, see the "Installing the Management Software" section on page 2-19.

Downgrading to Release 1.3(x) Versions

You can downgrade to Release 1.3(x) from a newer release, but you must first remove the old version. For example, you can downgrade from:

- Fabric Manager Release 3.x to version 1.x
- Fabric Manager Release 2.x to version 1.x

To downgrade, follow these steps:

Step 1 Close all instances of Fabric Manager Client or Device Manager on your workstation.

Step 2 Choose **Start > Programs > Cisco MDS 9000 > Uninstall** to uninstall Fabric Manager on Windows. Type **/usr/local/cisco_mds9000/uninstall.sh** or **$HOME/cisco_mds9000/uninstall.sh** to uninstall Fabric Manager on UNIX, depending on where Fabric Manager was installed.

Step 3 Enter the IP address or host name of the supervisor module with the lower version in the Address or Location field of your browser.

Step 4 Click the desired installation link (**Fabric Manager**, **Device Manager**, or **Fabric Manager Web Services and Performance Manager**).

Step 5 Select an installation folder for Fabric Manager on your workstation.

Unless you specify a different directory on a Windows PC, the version 1.3(x) software is installed in the default location of .\Documents and Settings\USER_ID\.cisco_mds9000. A Cisco MDS program group is created under Start > Programs. On a UNIX (Solaris or Linux) machine, the installation path name is /usr/local/.cisco_mds9000 or $HOME/.cisco_mds9000, depending on the permissions of the user doing the installation.

Launching the Management Software

To launch the Fabric Manager (Fabric View) or Device Manager (Device View and Summary View), follow these steps:

Step 1 Double-click either the **Fabric Manager** icon or the **Device Manager** icon on your desktop, or select the option from the Windows Start menu.

If you started Fabric Manager, the Fabric Manager Server loads You see a login screen for Fabric Manager or Device Manager. (You briefly see a command-line window.)

Step 2 Click **Options** to expand the login screen if necessary to select the seed switch and SNMP configuration.

Step 3 Enter the IP address or device name in the Device Name(s) field, or select an IP address from the list of previously accessed devices from the drop-down menu to the right of the Device Name(s) field.

Step 4 Leave the **SNMPv3** check box checked to select SNMP version 3. Otherwise, uncheck the check box to use SNMP version 2.

Note The default authentication digest used for storing user names and passwords is MD5. If you selected SHA instead, the related check box in the Fabric Manager initial login screen should be checked.

Step 5 Enter a user name and password.

Step 6 If the SNMPv3 Privacy option is enabled, enter the privacy password used for encrypting management traffic.

The Privacy option causes all management traffic to be encrypted while, with SNMPv3, user names and passwords are always encrypted.

Note You can create users with management traffic encryption (so a privacy password is required) or no management traffic encryption (no privacy password is required). Requiring a privacy password, and making it different from the authentication password, ensures stronger security but may cause AAA problems.

Step 7 Check the **Use SNMP Proxy** check box if you want Fabric Manager Client to communicate with Fabric Manager Server through a TCP-based proxy server. See the "Fabric Manager Server Proxy Services" section on page 2-15.

Note The **Accelerate Discovery** check box should remain checked for normal operation. Uncheck this only if you have changed switch IP addresses. You may experience problems with out of sync SAN IDs in Fabric Manager if you uncheck this check box.

Step 8 Optionally, select the **Local Interface** for Fabric Manager Client. Fabric Manager automatically detects the correct interface to use.

Step 9 Click **Open**.

You see either Fabric Manager or Device Manager.

Note

- When logging into Fabric Manager or Device Manager, the local SNMP database is checked first. If no user name entry is found, the AAA database is checked.

Note

- If you are using proxy services for Fabric Manager, and port 8080 is being used by another application, perform the following steps:

  - Open Internet Explorer
  - Select Tools > Internet Options > Connection > LAN settings
  - In the Proxy Server area, select Advance and add **localhost** or **ipaddress** under the exceptions

Note

- If you have an incomplete view of your fabric, rediscover the fabric with a user that has a network administrator or network operator role.

Integrating Cisco Fabric Manager with Other Management Tools

You can use Fabric Manager, Device Manager, and Performance Manager with other management tools. Here is a brief description of these tools. For more information on these tools and how they work together with the Cisco Fabric Manager management applications, see Chapter 62, "Troubleshooting Your Fabric,"

- Cisco Traffic Analyzer—Allows you to break down traffic by VSANs and protocols and to examine SCSI traffic at a logical unit number (LUN) level.
- Cisco Protocol Analyzer—Enables you to examine actual sequences of Fibre Channel frames easily using the Fibre Channel and SCSI decoders Cisco developed for Ethereal.
- Cisco Port Analyzer Adapter 2—Encapsulates SPAN traffic (both Fibre Channel control and data plane traffic) in an Ethernet header for transport to a Windows PC or workstation for analysis. Both the Cisco Traffic Analyzer and Cisco Protocol Analyzer require the PAA to transport MDS SPAN traffic to a Windows PC or workstation.

Running Fabric Manager Behind a Firewall

For Windows PCs running Fabric Manager, Device Manager, and Performance Manager behind a firewall, certain ports need to be available.

By default, Fabric Manager Client and Device Manager use the first available UDP port for receiving SNMP responses. The UDP SNMP trap local ports are 1162 for Fabric Manager, and 1163 or 1164 for Device Manager. Fabric Manager Client also opens TCP RMI port 9099. If Device Manager is opened from the Fabric Manager Client, it listens on the first available UDP port for Fabric Manager requests.

You can select the UDP port that Fabric Manager Client or Device Manager uses for SNMP responses by uncommenting the following statement:

- On a Windows desktop, uncomment the following in the FabricManager.bat or DeviceManager.bat file in the C:\Program Files\Cisco Systems\MDS9000\bin directory:

```
rem JVMARGS=%JVMARGS% -Dsnmp.localport=9001
```

- On a UNIX desktop, uncomment the following in the FabricManager.sh or DeviceManager.sh file in the $HOME/.cisco_mds9000/bin directory:

```
# JVMARGS=$JVMARGS -Dsnmp.localport=9001
```

Note UDP port 161 on the firewall must be open for incoming traffic. If the firewall blocks outgoing responses from snmp, then you can control which local ports DM or FM should open.

Fabric Manager Server proxy services uses a configurable TCP port (9189 by default) for SNMP communications between the Fabric Manager Client or Device Manager and Fabric Manager Server.

The Fabric Manager Server component requires two predictable TCP ports to be opened on the firewall for an incoming connection:

- java.rmi.registry.port = 9099
- java.rmi.server.remoteObjectPort = 9199

As long as these two ports are open, Fabric Manager Client can connect to the server. Other TCP ports connected to Fabric Manager Client are initiated by the server, which is behind the firewall.

The following table lists all ports used by Fabric Manager applications:

Communication Type	Port(s) Used
Used by All Applications	
SSH	Port 22 (TCP)
TELNET	Port 23 (TCP)
HTTP	Port 80 (TCP)
TFTP	Port 69 (UDP)
SYSLOG	Port 514 (UDP)
Used by Fabric Manager Server and Performance Manager	
SNMP_TRAP	Port 2162 (UDP)
SNMP	Picks a random free local port (UDP) – can be changed in server.properties.
Java RMI	Ports 9099, 9199 to 9299 (TCP)
Used by Fabric Manager Client	
Java RMI	Ports 9099, 9199 to 9299 (TCP)
SNMP	Picks a random free local port (UDP) or 9189 (TCP) if SNMP proxy is enabled – can be changed in **server.properties**.
Used by Device Manager	
SNMP_TRAP	Picks one available port in the range1163 to 1170 (UDP).
SNMP	Picks a random free local port (UDP) or 9189 (TCP) if SNMP proxy is enabled – can be changed in **server.properties**.

Uninstalling the Management Software

To uninstall the Fabric Manager applications on a Windows PC, follow these steps:

Step 1 Close all running instances of Fabric Manager and Device Manager.

Step 2 Select **Start > Programs > Cisco MDS 9000 > Uninstall** to run the uninstall.bat script.

You can also run the batch file (located in the C:\Program Files\Cisco Systems\MDS 9000 folder by default) directly from the command line.

Note For older installations, delete the .cisco_mds9000 folder. Manually delete all desktop icons and program menu items.

On a Windows PC, this folder is created under the Documents and Settings folder (for example, d:\Documents and Settings\Administrator\.cisco_mds9000 if you had installed it as user Administrator). On a UNIX machine, the default installation folder is /usr/bin.

Note You cannot downgrade from Fabric Manager Release 2.x to Fabric Manager Release 1.3(x). If you want to run Fabric Manager Release 1.3(x) on a PC that is running Fabric Manager Release 2.x, first uninstall Release 2.x then install Release 1.3. Fabric Manager does not work if you have both Release 2.x and Release 1.3 installed on the same PC.

To uninstall the Fabric Manager applications on a UNIX machine, follow these steps:

Step 1 For all releases starting with Release 2.x, run the shell script **$HOME/cisco_mds9000/Uninstall.sh** or **/usr/local/cisco_mds9000/uninstall.sh**, depending on where Fabric Manager was installed.

Step 2 For all releases starting with Release 1.3(1), run the shell script **$HOME/.cisco_mds9000/Uninstall.sh** or **/usr/local/.cisco_mds9000/uninstall.sh**, depending on where Fabric Manager was installed.

Step 3 For earlier installations, delete the $HOME/.cisco_mds9000 folder.

CHAPTER 3

Fabric Manager Server

Fabric Manager Server is a platform for advanced MDS monitoring, troubleshooting, and configuration capabilities. No additional software need be installed. The server capabilities are an integral part of the Cisco Fabric Manager software.

This chapter contains the following sections:

Fabric Manager Server Overview

Install Cisco Fabric Manager Server on a computer that you want to provide centralized MDS management services and performance monitoring. SNMP operations are used to efficiently collect fabric information. The Cisco Fabric Manager software, including the server components, requires about 20 MB of hard disk space on your workstation. Cisco Fabric Manager Server runs on Windows 2000, Windows 2003, Windows XP, Solaris 8.x or later, and Red Hat Linux.

Each computer configured as a Cisco Fabric Manager Server can monitor multiple Fibre Channel SAN fabrics. Up to 16 clients (by default) can connect to a single Cisco Fabric Manager Server concurrently. The Cisco Fabric Manager Clients can also connect directly to an MDS switch in fabrics that are not monitored by a Cisco Fabric Manager Server, which ensures you can manage any of your MDS devices from a single console.

Fabric Manager Server Features

Cisco Fabric Manager Server has the following features:

- Multiple fabric management— Fabric Manager Server monitors multiple physical fabrics under the same user interface. This facilitates managing redundant fabrics. A licensed Fabric Manager Server maintains up-to-date discovery information on all configured fabrics so device status and interconnections are immediately available when you open the Fabric Manager Client.

Note The unlicensed Cisco Fabric Manager can only monitor and configure one fabric at a time. You must use the Open menu to switch to a new fabric, which causes the application to stop monitoring the previous one and to rediscover the new fabric.

- Continuous health monitoring—MDS health is monitored continuously, so any events that occurred since the last time you opened the Fabric Manager Client are captured.
- Roaming user profiles—The licensed Fabric Manager Server uses the roaming user profile feature to store your preferences and topology map layouts on the server, so that your user interface will be consistent regardless of what computer you use to manage your storage networks.

Note You must have the same release of Fabric Manager Client and Fabric Manager Server.

Installing and Configuring Fabric Manager Server

Note Prior to running Fabric Manage Server, you should create a special Fabric Manager administrative user on each switch in the fabric or on a remote AAA server. Use this user to discover your fabric topology. See the "Best Practices for Discovering a Fabric" section on page 8-3.

To install Fabric Manager Server and set the initial configuration, follow these steps:

Step 1 Install Fabric Manager and Fabric Manager server on your workstation. See the "Installing Fabric Manager Server" section on page 3-3.

Step 2 Set the seed switch. See the "Setting the Seed Switch" section on page 3-4.

Step 3 Optionally, create flows and collections for Performance Manager to monitor your fabric. See the "Configuring Flows and Collections with Performance Manager" section on page 3-4.

Step 4 Set Fabric Manager Server to continuously monitor the fabric. See the "Fabric Manager Server Fabric Monitoring and Removal" section on page 3-6.

Step 5 Repeat Step 2 through Step 4 for each fabric that you want to manage through Fabric Manager Server.

Step 6 Install Web Services. See the "Installing Fabric Manager Web Services" section on page 3-6.

Step 7 Verify Performance Manager is collecting data. See the "Verifying Performance Manager Collections" section on page 3-6.

Installing Fabric Manager Server

When you install Fabric Manager, the basic version of the Fabric Manager Server (unlicensed) is installed with it. After you click the Fabric Manager icon, a dialog box opens and you can enter the IP address of a computer running the Fabric Manager Server component. If you do not see the Fabric Manager Server IP address text box, click **Options** to expand the list of configuration options. If the server component is running on your local machine, leave **localhost** in that field. If you try to run Fabric Manager without specifying a valid server, you are prompted to start the Fabric Manager Server locally.

On a Windows PC, you install the Fabric Manager Server as a service. This service can then be administered using Services in the Administrative Tools. The default setting for the Fabric Manager Server service is that the server is automatically started when the Windows PC is rebooted. You can change this behavior by modifying the properties in Services.

Unlicensed Versus Licensed Fabric Manager Server

When you install Fabric Manager, the basic unlicensed version of Fabric Manager Server is installed with it. To get the licensed features, such as Performance Manager, remote client support, and continuously monitored fabrics, you need to buy and install the Fabric Manager Server package.

However, trial versions of these licensed features are available. To enable the trial version of a feature, you run the feature as you would if you had purchased the license. You see a dialog box explaining that this is a demo version of the feature and that it is enabled for a limited time.

If you are evaluating one of these Fabric Manager Server features and want to stop the evaluation period for that feature, you can do that using Device Manager. See the "Fabric Manager Server Licensing" section on page 10-16.

Setting the Seed Switch

When you run Fabric Manager, you must select a switch for Fabric Manager to use to discover the fabric.

To set the seed switch, follow these steps:

Step 1 Double-click the **Fabric Manager Client** icon on your workstation. You see the Fabric Manager log in dialog box, as shown in Figure 3-1.

Figure 3-1 Starting Fabric Manager

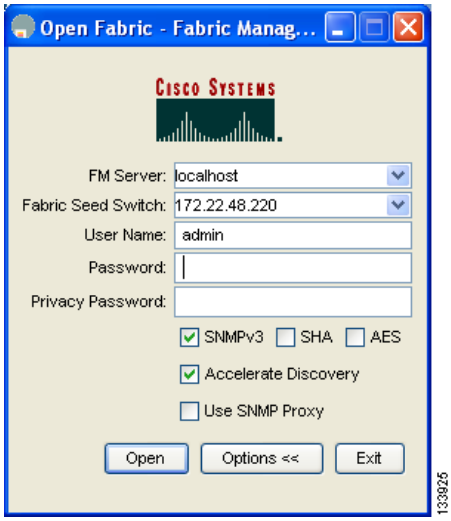

Step 2 Click **Options** if necessary to expand the optional settings in this dialog box.

Step 3 Set **FM Server** to the IP address where you installed the Fabric Manager Server or set it to **localhost** if you installed Fabric Manager Server on your local workstation.

Step 4 Set the Fabric Seed Switch to the MDS 9000 Family switch that you want Fabric Manager to use.

Step 5 Set the user name and password as required to start Fabric Manager Client.

Configuring Flows and Collections with Performance Manager

If you are managing your fabrics with Performance Manager, you need to set up an initial set of flows and collections on the fabric. See the "Historical Performance Monitoring" section on page 54-4 for a full description on Performance Manager features.

To create a flow in Performance Manager Configuration Wizard using Fabric Manager, follow these steps:

Step 1 Click **Performance > Create Flows** to launch the wizard.

Step 2 Choose the VSAN from which to create flows. Flows are defined per VSAN.

Step 3 Choose the **Type** radio button for the flow type you want to define.

Step 4 Check the **Clear old flows on modified switches** check box if you want to remove old flow data.

Step 5 Click **Next** to review the available flows for the chosen VSAN.

You see the screen shown in Figure 3-2.

Figure 3-2 Traffic Flows Found by the Flow Configuration Wizard

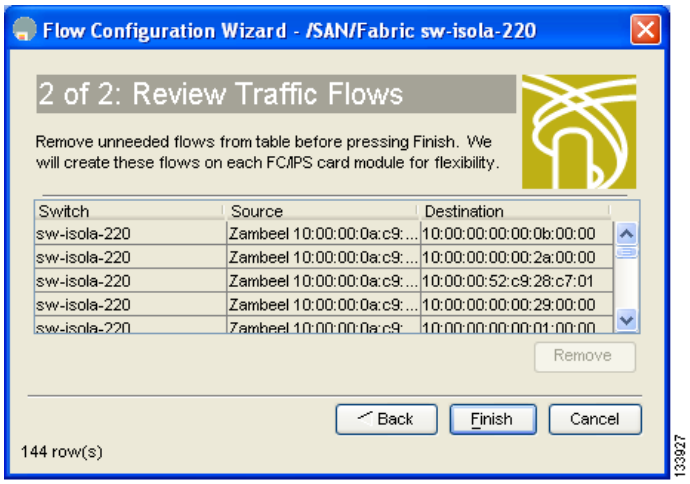

Step 6 Remove any flows you are not interested in.

Step 7 Click **Finish** to create the flow.

Using the Performance Manager Configuration Wizard

To create a collection using the Performance Manager Configuration Wizard in Fabric Manager Web Services, follow these steps:

Step 1 Click **Performance > Create Collection** to launch the Performance Manager Configuration Wizard.

Step 2 Choose the VSANs from which you want to collect data or choose **All** to collect statistics across all VSANs in the fabric.

Step 3 Check the **Type** check boxes for each type of link, flow, or SAN element that you want included in your collection.

Step 4 Check the check box to ignore flows with zero counter values, if appropriate.

Step 5 Enter the URL where Cisco Traffic Analyzer is located on your network, if appropriate.

Step 6 Click **Next** to review the collection specification data. Remove any links, flows, or SAN elements you are not interested in.

Step 7 Click **Next** to configure other collection options.

Step 8 Check the appropriate check boxes if you want to include errors and discards in your collection, and if you want to interpolate data for missing statistics.

Step 9 Check the **Send event if traffic exceeds threshold** check box if you want to configure threshold events as explained in the "Using Performance Thresholds" section on page 54-6.

Step 10 Click the **Use absolute values** radio button if you want absolute value thresholds, or click the **Baseline values over** radio button if you want baseline thresholds.

Step 11 Choose the time window for baseline calculations if baseline thresholds are configured.

Step 12 Choose the critical and warning threshold values as a percent of link capacity (for absolute value thresholds) or weighted average (for baseline thresholds).

Step 13 Click **Finish** to create the collection configuration file. You see a dialog box asking if you want to restart Performance Manager.

Step 14 Click **Yes** to restart Performance Manager to use this new configuration file, or click **No** to exit the Performance Manager Configuration Wizard without restarting Performance Manager. If you choose No, Performance Manager will not use the new configuration file until you restart it by choosing **Performance Manager > Collector > Restart.**

Installing Fabric Manager Web Services

You must install Fabric Manager Web Services to view Performance Manager reports through a web browser. To install Fabric Manager Web Services from the CD-ROM, see the "Installing Fabric Manager Web Services" section on page 6-3.

Verifying Performance Manager Collections

Once Performance Manager collections have been running for five or more minutes, you can verify that the collections are gathering data by choosing **Performance Manager > Reports** in Fabric Manager. You see the first few data points gathered in the graphs and tables.

Note Viewing reports requires installing Fabric Manager Web Services. See the "Installing Fabric Manager Web Services" section on page 6-3.

Fabric Manager Server Fabric Monitoring and Removal

Fabric Manager Server can continuously monitor a fabric, whether or not an instance of Fabric Manager (client) is monitoring that fabric. A continuously monitored fabric is automatically reloaded and monitored by Fabric Manager Server after the server starts. Fabrics that are monitored by Fabric Manager Server can have their data managed by Performance Manager. Both the Continuous Monitor feature and Performance Manager require the Fabric Manager Server license. However, you can "check out" these features without a license for a limited time.

Designating a Fabric for Continuous Monitoring

When you quit the Fabric Manager Client, you are prompted as to whether or not you would like to have Fabric Manager Server continuously monitor that fabric. Alternatively, you can use Fabric Manager Client to select a fabric to monitor.

To continuously monitor a fabric using Fabric Manager, follow these steps:

Step 1 Click **Server > Admin**.

You see a list of fabrics in the Server Admin dialog box.

Step 2 Check the **Monitor Continuously** check box next to the fabric(s) you want Fabric Manager Server to monitor.

Figure 3-3 Monitor Continuously Check Box Enabled

Step 3 Click **Apply**.

The Continuously Monitor feature requires the purchase of the Fabric Manager Server license package. If you have not purchased and installed this package, you see a pop-up window informing you that you are about to enable a demo license for this feature. Click **Yes** to enable the demo license.

Note When you are finished trying the licensed features, you can "check in" the feature by clicking the **Check In FM** button as described in the "Fabric Manager Server Licensing" section on page 10-16.

Step 4 Click **Close** to close the Server Admin dialog box.

Note If you will be collecting data on these fabrics using Performance Manager, you should now configure flows and define the data collections. These procedures are described in Chapter 53, "Network Monitoring."

Removing a Fabric from Monitoring

To remove a fabric from the Fabric Manager Server monitoring list using Fabric Manager, follow these steps:

Step 1 Click **Server > Admin**.

You see a list of fabrics in the Server Admin dialog box with the **Monitor Continuously** check box checked for one or more fabrics (see Figure 3-4).

Figure 3-4 Monitor Continuously Check Box Enabled

Step 2 Uncheck the **Continuously Monitor** check box next to the fabrics you no longer want Fabric Manager Server to monitor.

Step 3 Click **Apply**.

Step 4 Click **Close** to close the Server Admin dialog box.

Fabric Manager Server Properties File

The Fabric Manager Server properties file (MDS 9000\server.properties) contains a list of properties that determine how the Fabric Manager Server will function. You can edit this file with a text editor, or you can set the properties through the Fabric Manager Web Services GUI, under the Admin tab.

The server properties file contains these nine general sections:

- GENERAL—Contains the general settings for the server.
- SNMP SPECIFIC—Contains the settings for SNMP requests, responses, and traps.
- SNMP PROXY SERVER SPECIFIC—Contains the settings for SNMP proxy server configuration and TCP port designation.
- GLOBAL FABRIC—Contains the settings for fabrics, such as discovery and loading.
- CLIENT SESSION—Contains the settings for Fabric Manager Clients that can log into the server.
- EVENTS—Contains the settings for syslog messages.
- Performance Chart—Contains the settings for defining the end time to generate a Performance Manager chart.
- EMC CALL HOME—Contains the settings for the forwarding of traps as XML data using e-mail, according to EMC specifications.
- EVENT FORWARD SETUP—Contains the settings for forwarding events logged by Cisco Fabric Manager Server through e-mail.

The following are new or changed server properties for Fabric Manager Release 3.x:

SNMP Specific

- **snmp.preferTCP**—Specifies if SNMP over TCP should be used where possible. By default, this setting is **true**. The advantages of this setting are an increased buffer size and faster transfer rate. If your fabric has a long timeout period, it may prevent you from using SNMP (which may have a relatively shorter timeout period). If so, change this setting to **false** and restart Fabric Manager Server. UDP is used instead.

Note If it is set to false, the same choice must be set in Fabric Manager.

Performance Chart

- **pmchart.currenttime**—Specifies the end time to generate a Performance Manager chart. This should only be used for debugging purposes.

EMC Call Home

- **server.callhome.enable**—Enables or disables EMC Call Home. By default, it is disabled.
- **server.callhome.location**—Specifies the Location parameter.
- **server.callhome.fromEmail**—Specifies the 'From Email' list.
- **server.callhome.recipientEmail**—Specifies the 'recipientEmail' list.
- **server.callhome.smtphost**—Specifies the SMTP host address for outbound e-mail.
- **server.callhome.xmlDir**—Specifies the path to store the XML message files.
- **server.callhome.connectType**—Specifies the method to use to remotely connect to the server.
- **server.callhome.accessType**—Specifies the method to use to establish remote communication with the server.
- **server.callhome.version**—Specifies the version number of the connection type.
- **server.callhome.routerIp**—Specifies the public IP address of the RSC router.

Event Forwarding

- **server.forward.event.enable**—Enables or disables event forwarding.
- **server.forward.email.fromAddress**—Specifies the 'From Email' list.
- **server.forward.email.mailCC**—Specifies the 'CC Email' list.
- **server.forward.email.mailBCC**—Specifies the 'BCC Email' list.
- **server.forward.email.smtphost**—Specifies the SMTP host address for outbound e-mail.

For more information on setting the server properties, read the server.properties file or see the "Configuring Fabric Manager Server Preferences" section on page 6-42

Modifying Fabric Manager Server

Fabric Manager Release 2.1(2) or later allows you to modify certain Fabric Manager Server settings without stopping and starting the server. These settings include:

- Changing the Fabric Manager Server User Name and Password, page 3-10
- Changing the Polling Period and Fabric Rediscovery Time, page 3-10

Changing the Fabric Manager Server User Name and Password

You can modify the user name or password used to access a fabric from Fabric Manager Client without restarting Fabric Manager Server.

To change the user name or password used by Fabric Manager Server, follow these steps:

Step 1 Click **Server > Admin**.

You see the Server Admin dialog box displayed.

Step 2 Set the Name or Password for each fabric that you are monitoring with Fabric Manager Server.

You see the Server Admin dialog box shown in Figure 3-5.

Figure 3-5 Server Admin Dialog Box

Step 3 Click **Apply** to save these changes or click **Close** to exit the dialog box without saving any changes.

Changing the Polling Period and Fabric Rediscovery Time

Fabric Manager Server periodically polls the monitored fabrics and periodically rediscovers the full fabric at a default interval of five cycles.You can modify these settings from Fabric Manager Client without restarting Fabric Manager Server.

To change the polling period or full fabric rediscovery setting used by Fabric Manager Server using Fabric Manager, follow these steps:

Step 1 Select **Server > Admin**.

You see the Admin dialog box as shown in Figure 3-6.

Figure 3-6 Polling Interval Highlighted in the Server Admin Dialog Box

Step 2 For each fabric that you are monitoring with Fabric Manager Server, set the Polling Interval to determine how frequently Fabric Manager Server polls the fabric elements for status and statistics.

Step 3 For each fabric that you are monitoring with Fabric Manager Server, set the Rediscovery Cycles to determine how often Fabric Manager Server rediscovers the full fabric.

Step 4 Click **Apply** to save these changes or click **Close** to exit the dialog box without saving any changes.

Using Device Aliases or FC Aliases

You can change whether Fabric Manager uses FC aliases or global device aliases from Fabric Manager Client without restarting Fabric Manager Server.

To change whether Fabric Manager uses FC aliases or global device aliases using Fabric Manager, follow these steps:

Step 1 Click **Server > Admin**.

You see the Admin dialog box.

Step 2 For each fabric that you are monitoring with Fabric Manager Server, check the **Device Alias** check box to use global device aliases, or uncheck to use FC aliases (see Figure 3-7).

Figure 3-7 Device Alias Option Checked in the Server Admin Dialog box

Step 3 Click **Apply** to save these changes or click **Close** to exit the dialog box without saving any changes.

Saving Device Aliases to the Switch

If you choose to use global device aliases on Fabric Manager Server, these changes are not reflected on the local switch. The switch continues to use FC aliases until you save the device aliases to the switch.

To save global device aliases on a switch using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **End Devices** then select **Hosts** or **Storage** in the Physical Attributes pane. You see the end devices in the Information pane.

Step 2 For each device alias that you want the switch to recognize, highlight it, right-click the Device Alias icon, and select Save Selected Device Aliases.

CHAPTER 4

Fabric Manager Client

Cisco Fabric Manager Client is a java-based GUI application that provides access to the Fabric Manager applications from a remote workstation.

This chapter contains the following sections:

About Fabric Manager Client

Cisco Fabric Manager is a Java and SNMP-based network fabric and device management tool with a GUI that displays real-time views of your network fabric, including Cisco MDS 9000 Family and third-party switches, hosts, and storage devices.

In addition to complete configuration and status monitoring capabilities for Cisco MDS 9000 Family switches, Fabric Manager Client provides Fibre Channel troubleshooting tools. These health and configuration analysis tools use the MDS 9000 switch capabilities including Fibre Channel ping and traceroute.

Note You must use the same release of Fabric Manager Client and Fabric Manager Server.

Fabric Manager Advanced Mode

Advanced mode is enabled by default and provides the full suite of Fabric Manager features, including security, IVR, iSCSI, and FICON. Uncheck the **Advanced** check box in the upper right corner of the Fabric Manager Client to simplify the user interface. In this mode, you can access basic MDS 9000 features such as VSANs, zoning, and configuring interfaces.

Launching Fabric Manager Client

To launch Fabric Manager Client, follow these steps:

Step 1 Double-click the **Fabric Manager** icon on the Desktop.

Step 2 Follow the instructions described in the "Launching the Management Software" section on page 2-23 to launch Fabric Manager Client from your desktop.

To launch Fabric Manager Client from within a running instance of Fabric Manager, follow these steps:

Step 1 Choose **Open Fabric** from the File menu or click the **Open Switch Fabric** icon on the Fabric Manager toolbar.

You see the Open Fabric dialog box shown in Figure 4-1.

Figure 4-1 Open Fabric Dialog Box

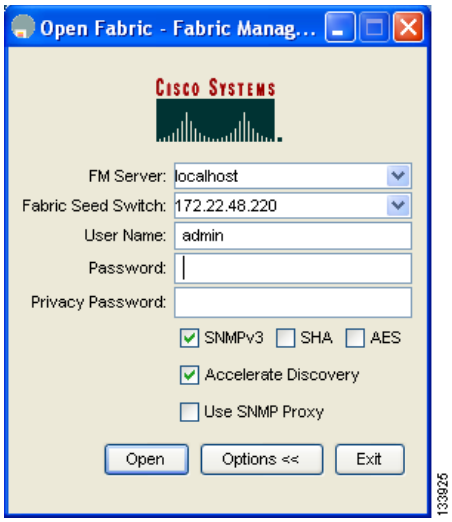

Step 2 Select the IP address of the fabric seed switch that you want to access. If you do not see the fabric seed switch in the pop-up window, click **Options** to expand the pop-up window options (see Figure 4-1).

Step 3 Click **Open** if the fabric that you want to open has the same user name and password as the fabric that you already have open. If the fabric that you want to open has a different user name and password, enter the user name and password and click **Open**. You are prompted for whether you want the previous fabric(s) to be monitored in the background. That fabric is then closed on Fabric Manager Client, and the new fabric is opened.

Fabric Manager displays the new fabric and adds a tab to the Fabric pane.

Step 4 Click each fabric's tab to view the fabric.

Note Changes made using Fabric Manager are applied to the running configuration of the switches that you are managing. If you have made changes to the configuration or performed an operation (such as activating zones), Fabric Manager prompts you to save your changes before you exit.

Fabric Manager Client Quick Tour

This section describes the Fabric Manager Client interface as shown in Figure 4-2.

Figure 4-2 Fabric Manager Main Window

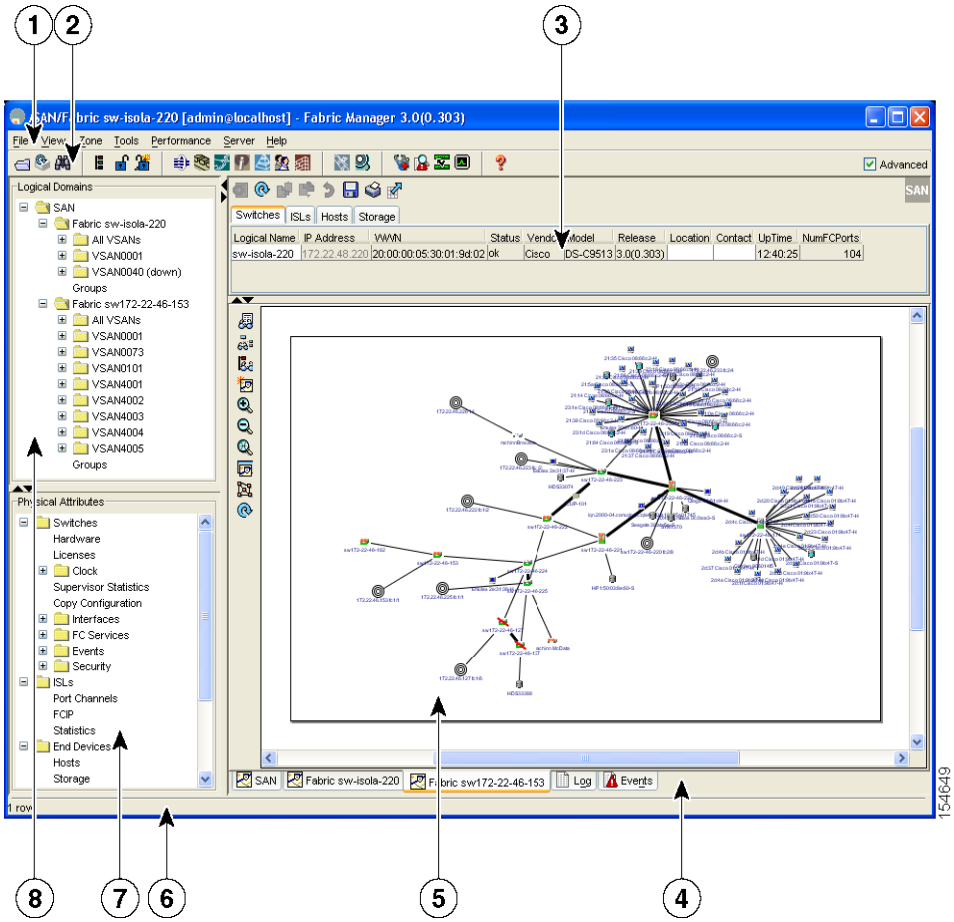

1	Menu bar—Provides access to options that are organized by menus.
2	Toolbar—Provides icons for direct access to the most commonly used options on the File, Tools, and Help menus.
3	Information pane—Displays information about whatever option is selected in the menu tree.
4	Status Bar (right side)—Shows the last entry displayed by the discovery process and the possible error message.
5	Fabric pane—Displays a map of the network fabric, including switches, hosts, and storage. It also provides tabs for displaying log and event data.
6	Status Bar (left side)—Shows short-term transient messages, such as the number of rows displayed in a table.
7	Logical Domains pane—Displays a tree of configured SAN, fabrics, VSANs, and zones, and provides access to user-defined groups.
8	Physical Attributes pane—Displays a tree of available configuration tasks depending on the fabric, VSAN, or zone selected previously. Lists the switches and end devices in the logical selection.

Note You can resize each pane by dragging the boundaries between each region or by clicking the **Minimize** or **Maximize** controls.

Menu Bar

The menu bar at the top of the Fabric Manager main window provides options for managing and troubleshooting the current fabric and for controlling the display of information on the Fabric pane. The menu bar provides the following menus:

- File—Opens a new fabric, rediscovers the current fabric, locates switches, sets preferences, prints the map, and clears (right-click on log) or exports the Fabric pane log.
- View—Changes the appearance of the map (these options are duplicated on the Fabric pane toolbar).
- Zone—Manages zones, zone sets, and inter-VSAN routing (IVR).
- Tools—Verifies and troubleshoots connectivity and configuration, as described in the "Fabric Manager Troubleshooting Tools" section on page 4-19.
- Performance—Runs and configures Performance Manager and Cisco Traffic Analyzer, and generates reports.
- Server—Runs administrative tasks on clients and fabrics. Provides Fabric Manager Server management and a **purge** command. Lists switches being managed.
- Help—Displays online help topics for specific dialog boxes in the Information pane.

Tool Bar

The Fabric Manager main toolbar provides icons for accessing the most commonly used menu bar options as shown inTable 4-1.

Table 4-1 Fabric Manager Client Main Toolbar

Icon	Description
	Opens switch fabric.
	Rediscovers current fabric.
	Finds in the map.
	Creates VSAN.
	Launches DPVM wizard.
	Edits full zone database.
	Launches IVR zone wizard.
	Launches PortChannel wizard.
	Launches FCIP wizard.
	Launches iSCSI wizard.
	Launches QoS wizard.
	Configures users and roles.

Table 4-1 Fabric Manager Client Main Toolbar (continued)

Icon	Description
	Launches IP-ACL wizard.
	Launches License Install wizard.
	Launches Software Install wizard.
	Performs switch health analysis.
	Performs fabric configuration analysis.
	Performs end-to-end connectivity analysis.
	Monitors ISL performance. Brings up real-time ISL performance information for all interfaces in the fabric, in the Information pane.
	Shows online help.

Logical Domains Pane

Use the Logical Domains pane to manage attributes for fabrics, VSANs, and zones, and to access user-defined groups. Right-click one of the folders in the tree and click a menu item from the pop-up menu. You see the appropriate configuration dialog box.

The default name for the fabric is the name, IP address, or WWN for the principal switch in VSAN 1. If VSAN 1 is segmented, the default name is chosen from a principal switch with the smallest WWN. In order, the fabric names you may see are as follows:

- Fabric <sysName>
- Fabric <ipAddress>
- Fabric <sWWN>

To change the fabric name from Fabric Manager, follow these steps:

Step 1 Select **Admin** from the Server menu.

You see the Server Admin dialog box.

Step 2 Double click the fabric name, and enter the new name of the fabric.

Step 3 Click **Apply** to change the name, and click **Close** to close the dialog box.

Filtering

Fabric Manager has a filtering mechanism that displays only the data that you are interested in. To filter, first select the fabric and VSAN from the Logical Domains pane. This action narrows the scope of what is displayed in the Fabric pane. Any information that does not belong to the selected items is dimmed. Also, any information that does not belong to the selected items is not displayed in the tables in the Information pane. As shown in Figure 4-3, the filter that you select is displayed at the top right of the Fabric Manager window.

To further narrow the scope, select attributes from the Physical Attributes pane. The Fabric Manager table, display, and filter criteria change accordingly.

Physical Attributes Pane

Use the Physical Attributes pane to display a tree of the options available for managing the switches in the currently selected fabric, VSAN, or zone.

To select an option, click a folder to display the options available and then click the option. You see the table with information for the selected option in the Information pane. The Physical Attributes pane provides the following main folders:

- Switches—Views and configures hardware, system, licensing, and configuration files.
- Interfaces—Views and configures FC physical, FC logical, Ethernet, SVC, and PortChannel interfaces.
- FC Services—Views and configures Fibre Channel network configurations.
- IP—Views and configures IP storage and IP services.
- Events—Views and configures events, alarms, thresholds, notifications, and informs.
- Security—Views and configures MDS management and FC-SP security.
- ISLs—Views and configures Inter-Switch Links.
- End Devices—Views and configures end devices.

Information Pane

Use the Information pane to display tables of information associated with the option selected from the menu tree in the Logical Domains or Physical Attributes panes. The Information pane toolbar provides buttons for performing one or more of the operations shown in Table 4-2.

Table 4-2 Information Pane Toolbar

Icon	Description
Apply Changes	Applies configuration changes.
Refresh Values	Refreshes table values.
Create Row	Opens the appropriate dialog box to make a new row in the table.
Delete Row	Deletes the currently highlighted rows from the table.
Copy/Ctrl+C	Copies data from one row to another.
Paste/Ctrl +V	Pastes the data from one row to another.
Undo Changes/Ctrl-Z	Undoes the most recent change.
Export	Exports and saves information to a file.
Print Table	Prints the contents of the Information pane.
Detach Table	Displays a non-editable copy of the table in the Information pane in its own window, which you can move around the screen.

Note After making changes, you must save the configuration or the changes will be lost when the device is restarted.

Note The buttons that appear on the toolbar vary according to the option that you select. They are activated or deactivated (dimmed) according to the field or other object that you select in the Information pane.

Detachable Tables

Detachable tables in Fabric Manager let you detach tables and move them to different areas on your desktop so that you can compare similar tables from different VSANs. You can keep informational tables open from one view while you examine a different area in Fabric Manager. To detach tables, click the **Detach Table** icon in the Information pane in Fabric Manager.

Fabric Pane

Use the Fabric pane to display the graphical representation of your fabric. Table 4-3 explains the graphics you may see displayed, depending on which devices you have in your fabric.

Table 4-3 Fabric Manager Graphics

Icon or Graphic	Description
	Director class MDS 9000.
	Non-director class MDS 9000.
	Generic Fibre Channel switch.
	Cisco SN5428.
	An orange line through a device indicates that the device is manageable but there are operational problems.
	An orange X through a device or link indicates that the device or ISL is not working properly.
	A red line through a device indicates that the device is not manageable.
	A red X through a device or link indicates that the device is down or that the ISL is down.
	Fibre Channel HBA (or enclosure).
	Fibre Channel target (or enclosure).

Table 4-3 Fabric Manager Graphics (continued)

Icon or Graphic	Description
	iSCSI host.
	Fibre Channel ISL and edge connection.
	Fibre Channel PortChannel.
	IP ISL and edge connection.
	IP PortChannel.
	Fibre Channel loop (storage).
	IP cloud (hosts). This icon is also used to represent a fabric when viewing a SAN (multiple fabrics) in the Fabric Manager Fabric pane.
	Any device, cloud, or loop with a box around it means that there are hidden links attached.

If a switch or director is grayed out, Fabric Manager can no longer communicate with it.

The bottom of the Fabric pane has the following tabs:

- Fabric—When displaying multiple fabrics, each fabric has its own tab. You can switch between fabrics by clicking on their respective tabs.
- Log—Displays messages that describe Fabric Manager operations, such as fabric discovery.
- Events—Displays information about the SNMP traps received by the management station. This includes combination events as detected by discovery and important traps like license, SNMP, and FICON.

When viewing large fabrics in the Fabric pane, it is helpful to do the following tasks:

- Turn off end device labels
- Collapse loops
- Collapse expanded multiple links (collapsed multiple links are shown as very thick single lines)
- Dim or hide portions of your fabric by VSAN

Note When a VSAN, zone, or zone member is selected in the VSAN tree, the map highlighting changes to identify the selected objects. To remove this highlighting, click the **Clear Highlight** button on the Fabric pane toolbar or choose **Clear Highlight** from the pop-up menu.

Context Menus

When you right-click an icon in the Fabric pane, you see a pop-up menu with options that vary depending on the type of icon selected. The various options available for different objects include the following:

- Open an instance of Device Manager for the selected switch.
- Open a CLI session for the selected switch.
- Copy the display name of the selected object.
- Execute a **ping** or **traceroute** command for the device.
- Show or hide end devices.
- View attributes.
- Quiesce and disable members for PortChannels.
- Set the trunking mode for an ISL.
- Create or add to a PortChannel for selected ISLs.

The Fabric pane has its own toolbar with options for saving, printing, and changing the appearance of the map. When you right-click the map, a pop-up menu appears that provides options (duplicated on the toolbar) for changing the appearance of the map.

Note You can launch web-based or non-web-based applications from the Fabric pane. To do this, you assign an IP address to the storage port or enclosure. Then right-click to bring up the pop-up menu, and select **Device Manager**.

Saving the Map

You can save the map in the Fabric Pane as an image, or as an editable Visio diagram. You can save the map with or without labels on the links. The created Visio diagram is editable and saved in two layers:

- The default layer includes all switches and links in the fabric.
- The end devices layer includes the end devices and can be turned off to remove end devices from the Visio diagram.

To save the map as a Visio diagram, select **Files > Export > Visio** and choose **Map** or **Map with link labels**. The saved Visio diagram retains the viewing options that you selected from the Fabric pane. For example, if you collapse multiple links in the map and export the links as a Visio diagram, the Visio diagram shows those multiple links as one solid link.

The Show Tech Support option from the Tools menu also supports saving the map as a Visio diagram.

Purging Down Elements

The Fabric pane allows you to refresh the map at any time by clicking the **Refresh Map** icon. The **Refresh Map** icon redraws the map but does not purge elements that are down. To purge down elements you can:

- Click **Server > Purge Down Elements**. This purges all down elements in the fabric.
- Right-click the **Fabric** pane and select **Purge Down Elements**.
- Right-click a down element and select **Purge**. This action purges only this element from the fabric.

Note If you select an element that is not down and purge it, that element will reappear on the next fabric discovery cycle.

Multiple Fabric Display

Fabric Manager can display multiple fabrics in the same pane (see Figure 4-3).

Figure 4-3 Fabric Manager's Multiple Fabric Display Window

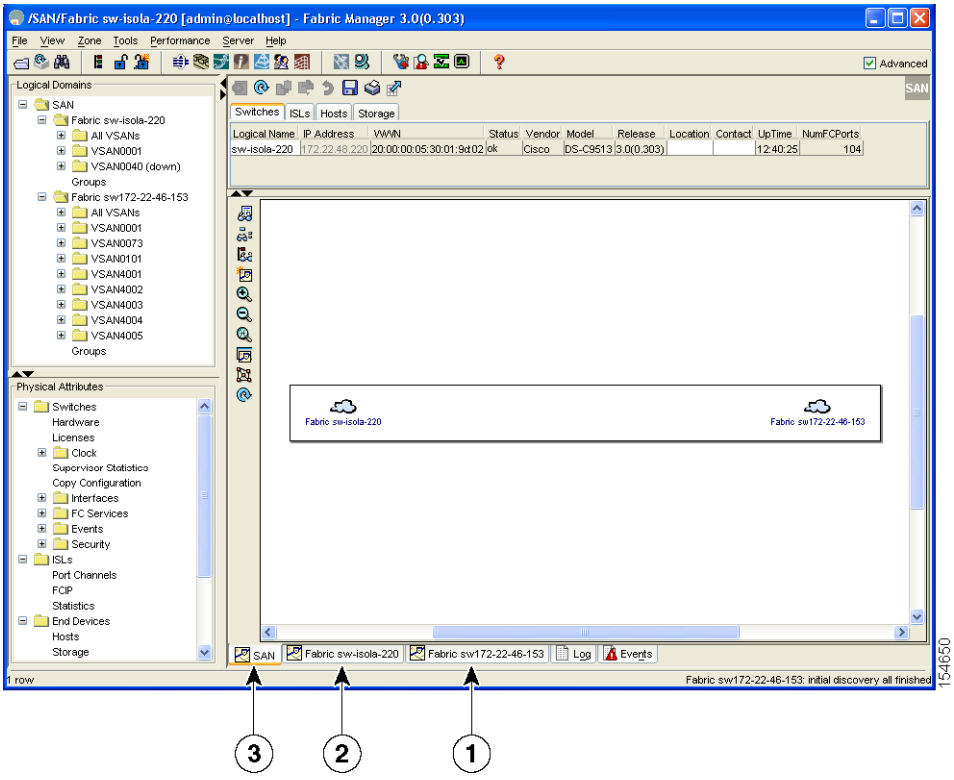

1	The Fabric view tab for fabric 172.23.46.152. When selected, the Fabric view displays fabric 172.23.46.152.
2	The Fabric view tab for fabric 172.23.46.153. When selected, the Fabric view displays fabric 172.23.46.153.
3	SAN tab (selected), showing two fabrics.

Note The same user name and password must be used to log into multiple fabrics.

The information for both fabrics is displayed; you do not need to select a seed switch. To see details of a fabric, select the tab for that fabric at the bottom of the Fabric pane, or double-click the **Cloud** icon for the fabric in the SAN tab.

Note Enclosure names should be unique. If the same enclosure name is used for each port, Fabric Manager shows a host/target enclosure connected to both fabrics. To fix this problem, you can either disable auto-creation or create unique enclosure names.

Filtering by Groups

You can filter the Fabric pane display by creating groups of switches or end ports. To create a group in Fabric Manager, follow these steps:

Step 1 Right-click a switch or end port in the Fabric pane map and select **Group > Create**.

You see the Edit User Defined Group dialog box shown in Figure 4-4.

Figure 4-4 Edit User Defined Group Dialog Box

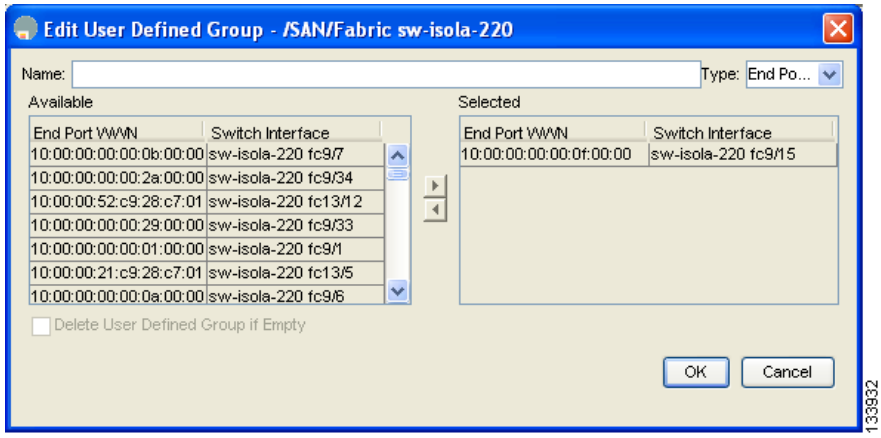

Step 2 Enter a group name in the Name field.

Step 3 Use the arrows to move additional switches or end ports from the Available column to the Selected column.

Step 4 Click **OK** to save the group.

To add a switch or end port to an existing group in Fabric Manager, follow these steps:

Step 1 Right-click a switch or end device and select **Group > Add To >** ***YourGroupName***.

You see the Edit User Defined Group dialog box, with your switch or end port added as shown in Figure 4-5.

Figure 4-5 Adding to a User Defined Group Dialog Box

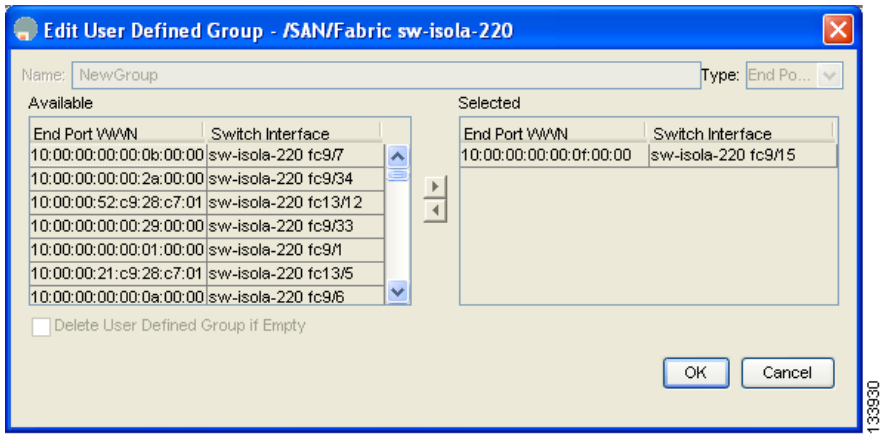

Step 2 Use the arrows to move additional switches or end ports from the Available column to the Selected column.

Step 3 Click **OK** to save the updated group.

To filter the display by a group you have created, follow these steps:

Step 1 Expand the **Groups** folder in the Logical Domains pane.

You see the list of groups that you have created as shown in Figure 4-6.

Figure 4-6 Group Highlighted in Fabric Pane Map

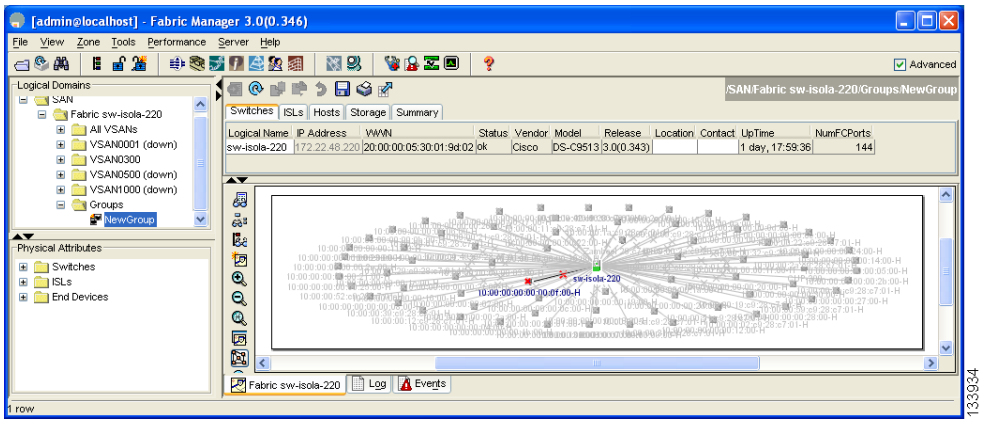

Step 2 Click the name of the group that you want to filter.

In the Fabric pane, the switches or end devices in your group are shown normally; all other switches and end devices are shown in gray.

Step 3 Click the **Groups** folder in the Logical Domains pane to return the display to normal.

Note User defined groups tables are filtered based on switches in the group except for switches where CFS-controlled features are enabled when all CFS member switches are displayed to avoid misconfigurations.

Status Bar

The status bar at the bottom of the Fabric Manager window shows the last entry displayed by the discovery process, and the possible error message on the right side. The status bar displays a message stating that something has changed in the fabric and a new discovery is needed. The status bar shows both short-term, transient messages (such as the number of rows displayed in the table) and long-term discovery issues.

Setting Fabric Manager Preferences

To set your preferences for the behavior of the Fabric Manager, choose **File > Preferences** from the Fabric Manager menu bar. You see the Preferences dialog box with the following tabs for setting different components of the application:

- General

- SNMP
- Map

The default General preferences for Fabric Manager are as follows:

- Show Device Name by—Displays the switches in the Fabric pane by IP address, DNS name, or logical name. The default setting for this value is Logical Name.
- Show WorldWideName (WWN) Vendor—Displays the world wide name vendor name in any table or listing displayed by Fabric Manager. Check the **Prepend Name** check box to display the name in front of the IP address of the switch. Check the **Replacing Vendor Bytes** check box to display the name instead of the IP address. The default is the Prepend Name option.
- Show End Device Using—Displays end devices in the Fabric pane using alias or pWWN alias. The default setting for this value is Alias.
- Show Shortened iSCSI Names—Displays the default setting for this value is OFF.
- Show Timestamps as Date/Time—Displays timestamps in the date/time format. If this preference is not checked, timestamps are displayed as elapsed time. The default setting is enabled (checked).
- Telnet Path—Displays the path for the telnet.exe file on your system. The default is telnet.exe, but you need to browse for the correct location.

Note If you browse for a path or enter a path and you have a space in the pathname (for example, c:\program files\telnet.exe), then the path will not work. To get the path to work, you must manually place quotes around it (for example, "c:\program files\telnet.exe**"**).

- Use Secure Shell instead of Telnet—Specifies whether to use SSH or Telnet when using the CLI to communicate with the switch. If enabled, you must specify the path to your SSH application. The default setting is disabled.
- Confirm Deletion—Displays a confirmation pop-up window when you delete part of your configuration using Fabric Manager. The default setting is enabled (checked).
- Export Tables with Format—Specifies the type of file that is created when you export a table using Device Manager. The options are tab-delimited or XML. The default setting is Tab-Delimited.
- Show CFS Warnings—Shows warning messages if CFS is not enabled on all switches for a selected feature.

The default SNMP preferences for Fabric Manager are as follows:

- Retry request 1 time(s) after 5 sec timeout—You can set the retry value to 0-5, and the timeout value to 3-30.
- Trace SNMP packets in Log—The default setting for this value is OFF.
- Enable Audible Alert when Event Received—The default setting for this value is OFF.

The default Map preferences for Fabric Manager are as follows:

- Display Unselected VSAN Members—Displays the unselected VSAN members in the Fabric pane. The default setting for this value is ON.
- Display End Devices—Displays the fabric's end devices in the Fabric pane. The default setting for this value is ON.
- Display End Device Labels—Displays the fabric's end device labels in the Fabric pane. The default setting for this value is ON.
- Expand Loops—Displays the loops in the fabric as individual connections in the Fabric pane. The default setting for this value is OFF.

- Expand Multiple Links—Displays multiple links in the Fabric pane as separate lines instead of one thick line. The default setting for this value is ON.
- Open New Device Manager Each Time—Opens a new instance of Device Manager each time that you invoke it from a switch in your fabric. The default value is OFF, which means that only one instance of Device Manager is open at a time.
- Select Switch or Link from Table—Allows you to select a switch or link in the Fabric pane by clicking the switch or link in a table in the Information pane. The default setting for this value is disabled (unchecked), which means clicking a switch or link in the table does not change the switch or link selection in the Fabric pane.
- Layout New Devices Automatically—Automatically places new devices in the Fabric pane in an optimal configuration. The default setting for this value is OFF. In this mode, when you add a new device, you must manually reposition it if the initial position does not suit your needs.
- Use Quick Layout when Switch has 30 or more End Devices—Displays the default setting for this value (30). You can enter any number in this field. Enter **0** to disable Quick Layout.
- Override Preferences for Non-default Layout—Displays the default setting for this value (ON).
- Automatically Save Layout—If this option is enabled, any changes in the layout are automatically saved. The default setting for this value is ON.
- Detach Overview Window—Allows you to easily center the Fabric pane on the area of the fabric that you want to see. (This feature is useful for large fabrics that cannot be displayed entirely within the Fabric pane.) Bring up the overview window by clicking the **Show/Hide Overview Window** button. It overlays the fabric window and remains there until you click the **Show/Hide Overview Window** button again. If you enable this preference, you can detach the overview window and move it to one side while you access the Fabric pane. The default setting for this value is disabled (unchecked).

Network Fabric Discovery

Cisco Fabric Manager collects information about the fabric topology through SNMP queries to the switches that are connected to Fabric Manager. The switch replies after having discovered all devices connected to the fabric by using the information from its FSPF technology database and the Name Server database and collected using the Fabric Configuration Server's request/response mechanisms that are defined by the FC-GS-3/4 standard. When you start Fabric Manager, you enter the IP address (or host name) of a seed switch for discovery.

After you start Fabric Manager and the discovery completes, Fabric Manager presents you with a view of your network fabric, including all discovered switches, hosts, and storage devices.

Modifying the Device Grouping

Because not all devices are capable of responding to FC-GS-3 requests, different ports of a single server or storage subsystem may be displayed as individual end devices on the Fabric Manager map.

To group end devices in a single enclosure to have them represented by a single icon on the map, Fabric Manager, follow these steps:

Step 1 Expand **End Devices** and then select **Storage** or **Hosts** in the Physical Attributes pane.

You see the end devices displayed in the Information pane.

Step 2 Click one of the devices in the Fabric pane, or click the **Enclosures** tab of the Information pane, then click the device name (in the Name field) you want to include in the enclosure.

Step 3 Enter a name to identify the new enclosure in the Fabric pane map.

Step 4 Click once on the device name in the Name field. To select more than one name, press the **Shift** key and click each of the other names.

Step 5 Press **Ctrl-C** to copy the selected name(s).

Step 6 Press **Ctrl-V** to paste the device name into the Name field.

Note To remove devices from an enclosure, triple click the device name and press **Delete**. To remove an enclosure, repeat this step for each device in the enclosure.

Using Alias Names as Enclosures

To create an enclosure that uses the alias name as the name of the enclosure using Fabric Manager, follow these steps:

Step 1 Expand End Devices and select **Hosts** or **Storage** from the Physical Attributes pane.

You see the list of devices in the Information pane. The NxPorts tab is the default.

Step 2 Right-click the enclosure names that you want to convert to alias names and select **Alias > Enclosure** as shown in Figure 4-7.

Figure 4-7 Alias Enclosure

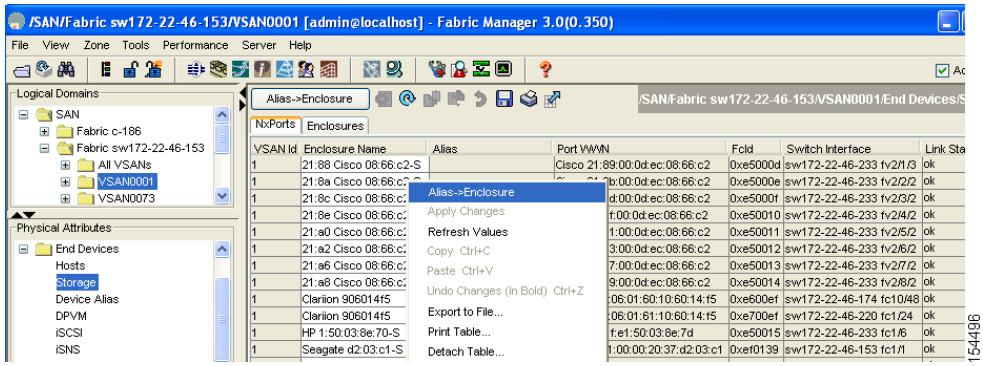

Step 3 Click the **Apply Changes** icon to save these changes or click **Undo Changes** to discard any unsaved changes.

Controlling Administrator Access with Users and Roles

Cisco MDS 9000 Family switches support role-based management access whether using the CLI or Cisco Fabric Manager. This lets you assign specific management privileges to particular roles and then assign one or more users to each role.

Cisco Fabric Manager uses SNMPv3 to establish role-based management access. After completing the setup routine, a single role, user name, and password are established. The role assigned to this user allows the highest level of privileges, which includes creating users and roles. Use the Cisco Fabric Manager to create roles and users and to assign passwords as required for secure management access in your network.

Using Fabric Manager Wizards

Fabric Manager Client provides the following wizards to facilitate common configuration tasks:

- VSAN—Creates VSANs on multiple switches in the fabric and sets VSAN attributes including interop mode, load balancing, and FICON.
- Zone Edit Tool—Creates zone sets, zones, and aliases. Adds members to zones and edits the zone database.
- IVR Zone—Creates IVR zone sets, zones, and aliases. Enables IVR NAT and auto-topology. Adds members to IVR zones, and edits the IVR zone database.
- PortChannel—Creates PortChannels from selected ISLs either manually or automatically. Sets PortChannel attributes such as channel ID and trunking mode.
- FCIP—Creates FCIP links between Gigabit Ethernet ports. Enables Fibre Channel write acceleration and IP compression.
- DPVM—Establishes dynamic port VSAN membership, enables autolearning, and activates the DPVM database.
- iSCSI—Creates zones for iSCSI initiators and adds a VSAN to a target-allowed VSAN list.
- QoS—Sets QoS attributes for zones in the selected VSAN.
- IP ACL—Creates ordered IP access control lists and distributes to selected switches in the fabric.
- License Install—Facilitates download and installation of licenses in selected switches in the fabric.
- Software Install—Verifies image compatibility and installs software images on selected switches in the fabric.

Fabric Manager Troubleshooting Tools

Fabric Manager has several troubleshooting tools available from the toolbar or Tools menu. Procedures for using these tools are described in Chapter 62, “Troubleshooting Your Fabric.” This section provides a brief description of each tool:

- Zone Merge Analysis—The zone merge analysis tool (available from the Zone menu) enables you to determine if zones will merge successfully when two Cisco MDS switches are interconnected. If the interconnected switch ports allow VSANs with identical names or contain zones with identical names, then Fabric Manager verifies that the zones contain identical members. The merge analysis tool can be run before attempting a merge or after fabrics are interconnected to determine zone merge failure causes.

- End-to-End Connectivity—Fabric Manager's end-to-end connectivity analysis tool uses FC Ping to verify interconnections between Cisco MDS switches and end-device (HBAs and storage devices) in a particular VSAN. In addition to basic connectivity, Fabric Manager can optionally verify the following:
  - Paths are redundant.
  - Zones contain at least two members.

End devices are connected to a manageable switch (have a currently active in-band or out-of-band management path.)

- Switch Health Analysis—You can run an in-depth switch health analysis with Fabric Manager. It verifies the status of all critical Cisco MDS switches, modules, ports, and Fibre Channel services. Over 40 conditions are checked. This tool provides a very fast, simple, and thorough way to assess Cisco MDS switch health.
- Fabric Configuration Analysis—Fabric Manager includes a fabric configuration analysis tool. It compares the configurations of all Cisco MDS switches in a fabric to a reference switch or a policy file. You can define what functions to check and what type of checks to perform. The analysis can look for mismatched values, and missing or extra values. If all configuration checking is performed for all functions, over 200 checks are performed for each Cisco MDS switch.

After the analysis is run, the results are displayed with details about the issues that were discovered. You can automatically resolve configuration differences by selecting them and clicking the **Resolve** button. Fabric Manager automatically changes the configuration to match the reference switch or policy file.

CHAPTER 5

Device Manager

This chapter contains descriptions of, and instructions for using, the Cisco MDS 9000 Device Manager. This chapter contains the following sections:

Device Manager Overview

Device Manager provides a graphic representation of a Cisco MDS 9000 Family switch chassis, including the installed switching modules, the supervisor modules, the status of each port within each module, the power supplies, and the fan assemblies.

The tables in the Fabric Manager Information pane basically correspond to the dialog boxes that appear in Device Manager. However, while Fabric Manager tables show values for one or more switches, a Device Manager dialog box shows values for a single switch. Also, Device Manager provides more detailed information for verifying or troubleshooting device-specific configuration than Fabric Manager.

Device Manager Features

Device Manager provides two views, Device View and Summary View. Use Summary View to monitor interfaces on the switch. Use Device View to perform switch-level configurations including the following:

- Manage ports, PortChannels, and trunking.
- Manage SNMPv3 security access to switches.
- Manage CLI security access to the switch.
- Manage alarms, events, and notifications.
- Save and copy configuration files and software image.
- View hardware configuration.
- View chassis, module, port status, and statistics.

Launching Device Manager

You can launch Device Manager two ways.

To launch Device Manager from your desktop, double-click the **Device Manager** icon and follow the instructions described in the "Launching the Management Software" section on page 2-23.

Launch Device Manager from Fabric Manager, any of three ways:

- Right-click the switch you want to manage on the Fabric pane map and choose **Device Manager** from the menu that appears.
- Double-click a switch in the Fabric pane map.
- Select a switch in the Fabric pane map and choose **Tools > Device Manager**.

Using Device Manager

This section describes the Device Manager interface, as shown in Figure 5-1.

Figure 5-1 Device Manager, Device Tab

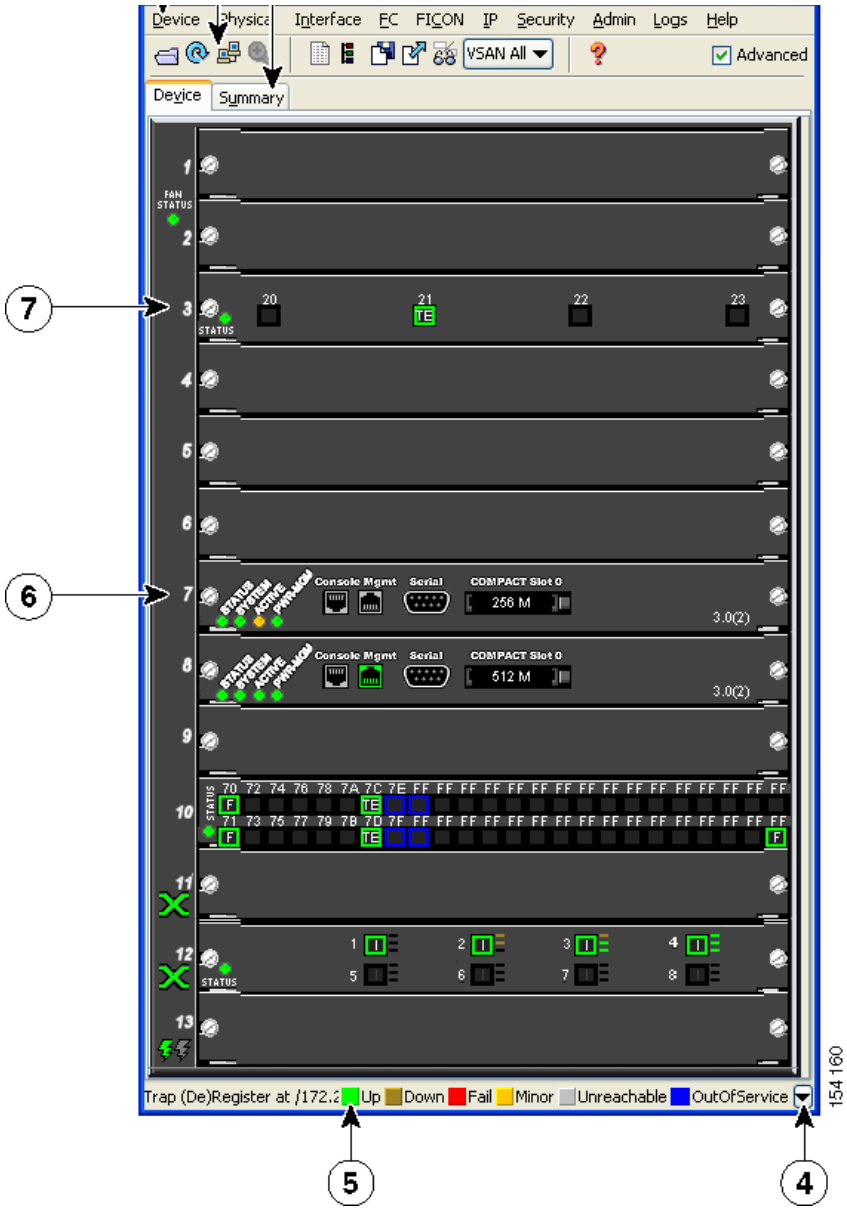

1	Menu bar	**5**	Status
2	Toolbar	**6**	Supervisor modules
3	Tabs	**7**	Switching or services modules
4	Legend		

Menu Bar

The menu bar at the top of the Device Manager main window provides options for managing and troubleshooting a single switch. The menu bar provides the following options:

- Device—Opens an instance of Device Manager, sets management preferences, sets the page layout, opens a Telnet/SSH session with the current switch, exports a device image, and closes the Device Manager application.
- Physical—Allows you to view and manage inventory, modules, temperature sensors, power supplies, fans, and the entire system.
- Interface—Allows you to configure and manage PortChannels, as well as Fibre Channel, Ethernet, iSCSI, and FICON ports. Also provides diagnostic, management and monitoring capabilities, as well as SPAN and port tracking.
- FC—Allows you to configure and manage VSAN, domain, and name server characteristics. Also provides advanced configuration capabilities.
- FICON—Allows you to configure and manage FICON VSANs, configure RLIR ERL information, swap selected FICON ports, and view FICON port numbers.
- IP—Allows you to configure and manage the following types of information: FCIP, iSCSI, iSNS, routes, VRRP, and CDP.
- Security—Allows you to configure and manage FCSP, port security, iSCSI security, SNMP security, common roles, SSH, AAA, and IP ACLs.
- Admin—Allows you to save, copy, edit, and erase the switch configuration, monitor events, manipulate Flash files, manage licenses, configure NTP, use CFS, and reset the switch. Also enables you to use the **show tech support, show cores**, and **show image** commands.
- Logs—Shows the various logs: message, hardware, events, and accounting. Also displays FICON link incidents, and allows you to configure the syslog setup.
- Help—Displays online help topics for specific dialog boxes in the Information pane.

Toolbar Icons

The Device Manager toolbar provides quick access to many Device Manager features. Once the icon is selected, a dialog box may open that allows configuration of the feature. The toolbar provides the main Device and Summary View icons as shown in Table 5-1.

Table 5-1 Device Manager Main Toolbar

Icon		Description
	Open Device	Opens the Device Manager view for another switch, with the option to open this view in a separate window.
	Refresh Display	Communicates with the switch and displays the information in the Device Manager view.
	Command-Line Interface	Opens a separate CLI command window to the switch.

Table 5-1 Device Manager Main Toolbar (continued)

Icon	Description
 Configure Selected	Opens a configuration dialog box for the selected component (line card or port).
 SysLog	Opens a window that lists the latest system messages that occurred on the switch.
 Threshold Manager	Opens the Threshold Manager dialog box that provides statistical monitoring and event reporting for the switch.
 VSANs	Opens the VSAN dialog box that provides VSAN configuration for the switch.
 SNMP Users and Roles	Opens the SNMP configuration dialog box for SNMP users and roles.
 Save Configuration	Saves the current running configuration to the startup configuration.
 Copy	Copies configuration file between server and switch
 Toggle FICON/Interface Port Labels	Toggles the FICON and interface port labels.
 Select VSAN	Filters the port display to show only those ports belonging to the selected VSAN.
 Help	Accesses online help for Device Manager.

Dialog Boxes

If a toolbar icon is selected, a dialog box may open that allows configuration of the selected feature. The dialog box may include table manipulation icons. See the "Information Pane" section on page 4-7 for descriptions of these icons.

Tabs

Click the **Device** tab on the Device Manager main window to see a graphical representation of the switch chassis and components.

Click the **Summary** tab on the Device Manager main window to see a summary of active interfaces on a single switch, as well as Fibre Channel and IP neighbor devices. The Summary View also displays port speed, link utilization, and other traffic statistics. There are two buttons in the upper left corner of the Summary View tab used to monitor traffic. To monitor traffic for selected objects, click the **Monitor Selected Interface Traffic Util%** button. To display detailed statistics for selected objects, click the **Monitor Selected Interface Traffic Details** button. You can set the poll interval, the type or Rx/Tx display, and the thresholds.

Legend

The legend at the bottom right of the Device Manager indicates port status, as follows:

Colors

- Green—The port is up.
- Brown—The port is administratively down.
- Red—The port is down or has failed.
- Amber—The port has a minor fault condition.
- Gray—The port is unreachable.

Labels

- X—Link failure
- E—ISL
- TE—Multi-VSAN ISL
- F—Host/storage
- FL—F loop
- I— iSCSI
- SD—SPAN destination
- CH—Channel
- CU—Control unit

Note For a detailed table describing the legend, see the "What does a red/orange/dotted line through the switch mean?" section on page 63-14.

Supervisor and Switching Modules

In the Device View, you can right-click an object and get information on it, or configure it. If you right-click a module, the menu shows the module number and gives you the option to configure or reset the module. If you right-click a port, the menu shows the port number and gives you the option to configure, monitor, enable/disable, set beacon mode, or perform diagnostics on the port.

Tip You can select multiple ports in Device Manager and apply options to all the selected ports at one time. Either select the ports by clicking the mouse and dragging it around them, or hold down the **Control** key and click each port.

To enable or disable a port, right-click the port and click **Enable** or **Disable** from the pop-up menu. To enable or disable multiple ports, drag the mouse to select the ports and then right-click the selected ports. Then click **Enable** or **Disable** from the pop-up menu.

To manage trunking on one or more ports, right-click the ports and click **Configure**. In the dialog box that appears, right-click the current value in the Trunk column and click **nonTrunk**, **trunk**, or **auto** from the pull-down list.

To create PortChannels using Device Manager, click **PortChannels** from the Interface menu. For detailed instructions, see Chapter 21, "Configuring PortChannels." You can also use Fabric Manager to conveniently create a PortChannel.

Note To create a PortChannel, all the ports on both ends of the link must have the same port speed, trunking type, and administrative state.

Context Menus

Context menus are available in both Device Manager views by right-clicking a device or table.

From Device View:

- Device—Right-click a system, module, or power supply to bring up a menu that gives you the option to configure or reset the device.
- Port— Right-click a port to bring up a menu that shows you the number of the port you have clicked, and to give you the option to configure, monitor, enable, disable, set beacon mode, or perform diagnostics on the port.

From Summary View:

- Table— Right-click the table header to show a list of which columns to display in that table: Interface, Description, VSANs, Mode, Connected To, Speed (Gb), Rx, Tx, Errors, Discards, and Log. Click the Description field to bring up the appropriate configuration dialog box for the port type.

Setting Device Manager Preferences

To set your preferences for the behavior of the Device Manager application, choose **Device > Preferences** from the Device menu. You can set the following preferences:

- Retry Requests x Time(s) After x sec Timeout—Allows you to set the retry request values. The default settings are 1 time after a 5-second timeout.
- Enable Status Polling Every x secs—Allows you to set the status polling value. The default setting is enabled (checked) with a time of 40 seconds.
- Trace SNMP Packets in Message Log—Allows you to set whether Device Manager traces SNMP packets and logs the trace. The default setting is disabled (unchecked).
- Register for Events After Open, Listen on Port 1163—Allows you to register this switch so that events are logged once you open Device Manager. The default setting is enabled (checked).
- Confirm Deletion—Displays a pop-up confirmation when you delete part of your configuration using Device Manager. The default setting is enabled (checked).
- Show WorldWideName (WWN) Vendor—Displays the world wide name vendor name in any table or listing displayed by Device Manager. If **Prepend** is checked, the name is displayed in front of the IP address of the switch. If **Replace** is checked, the name is displayed instead of the IP address. The default setting is enabled (checked) with the **Prepend** option.
- Show Timestamps as Date/Time—Displays timestamps in the date/time format. If this preference is not checked, timestamps are displayed as elapsed time. The default setting is enabled (checked).
- Telnet Path—Sets the path for the telnet.exe file on your system. The default is telnet.exe, but you need to browse for the correct location.

Note If you browse for a path or enter a path and you have a space in the pathname (for example, c:\program files\telnet.exe, then the path will not work. To get the path to work, manually place quotes around it (for example, "c:\program files\telnet.exe").

- Use Secure Shell Instead of Telnet—Specifies whether to use SSH or Telnet when using the CLI to communicate with the switch. If enabled, you must specify the path to your SSH application. The default setting is disabled.
- CLI Session Timeout x secs (0= disable)—Specifies the timeout interval for a CLI session. Enter 0 to disable (no timeout value). The default setting is 30 seconds.
- Show Tooltips in Physical View—Determines whether tooltips are displayed in Physical (Device) View. The default setting is enabled (checked).
- Label Physical View Ports With:—Specifies the type of label to assign to the ports when you are in Physical (Device) View. The options are FICON and Interface. The default setting is Interface.
- Export Table—Specifies the type of file that is created when you export a table using Device Manager. The options are Tab-Delimited or XML. The default setting is Tab-Delimited.

CHAPTER 6

Fabric Manager Web Services

With Fabric Manager Web Services you can monitor Cisco MDS switch events, performance, and inventory from a remote location using a web browser. This chapter contains the following sections:

Fabric Manager Web Services Overview

Using Fabric Manager Web Services, you can monitor MDS switch events, performance, and inventory, and perform minor administrative tasks.

Fabric Manager Web Services provides the following features:

- Summary and drill down reports—The Performance Manager summary report provides a high-level view of your network performance. These reports list the average and peak throughput and provides hot-links to additional performance graphs and tables with additional statistics. Both tabular and graphical reports are available for all interconnections monitored by Performance Manager. Performance Manager also analyzes daily, weekly, monthly and yearly trends. These reports are only available if you create a collection using Performance Manager and start the collector. See the “Historical Performance Monitoring” section on page 54-4.
- Zero maintenance database for statistics storage—No maintenance is required to maintain Performance Manager’s round-robin database, because its size does not grow over time. At prescribed intervals the oldest samples are averaged (rolled-up) and saved. A full two days of raw samples are saved for maximum resolution. Gradually the resolution is reduced as groups of the oldest samples are rolled up together.

Fabric Manager Web Services displays in a web browser shown in Figure 6-1.

Figure 6-1 Fabric Manager Web Services.

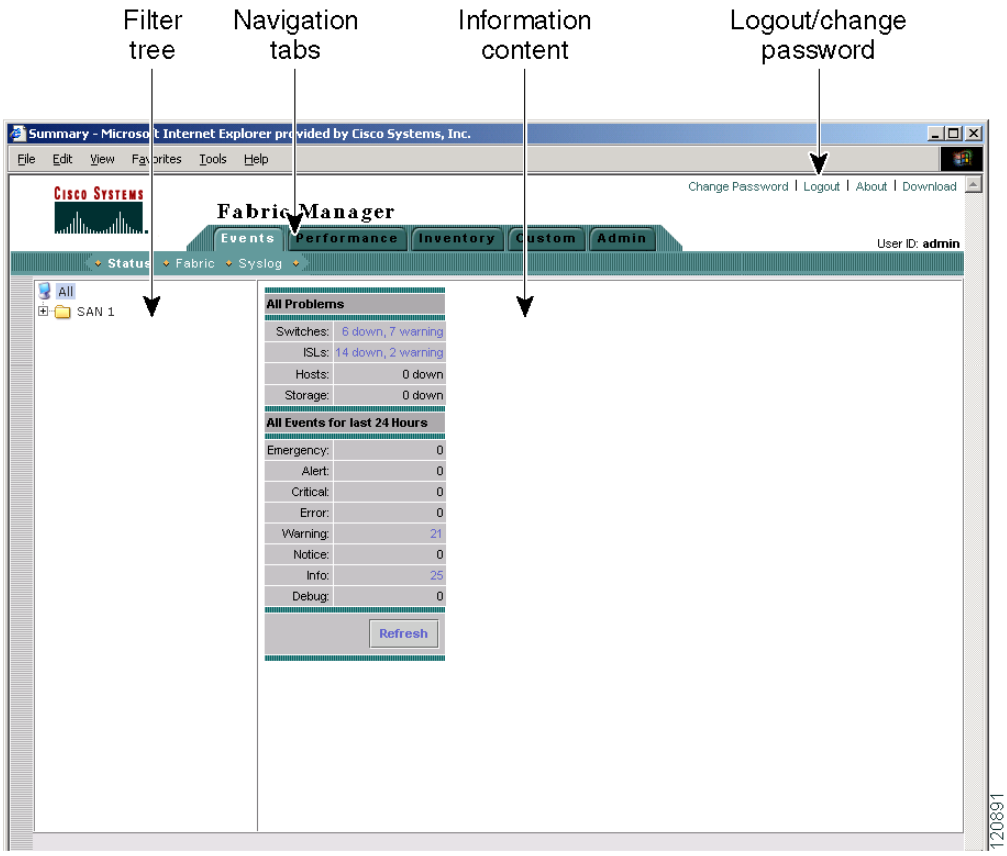

Installing Fabric Manager Web Services

If you are installing the Fabric Manager Web Services software for the first time, or if you want to update or reinstall the software, you access the supervisor module of the switch using a web browser. Install Fabric Manager Web Services on the same workstation where you installed Fabric Manager Server.

You must install Fabric Manager Web Services to view Performance Manager reports through a web browser.

For switches running Cisco MDS 9000 FabricWare, you need to install the Fabric Manager Web Services software from the CD-ROM included with your switch, or download Fabric Manager from Cisco.com.

To install Fabric Manager Web Services from the CD-ROM, navigate to the Fabric Manager installation notes and follow the directions.

To download the software from Cisco.com (requires a valid user name and password), go to the following website:

http://cisco.com/cgi-bin/tablebuild.pl/mds-fm

To download and install the software on your workstation, follow these steps:

Step 1 Optionally, enter the IP address or host name of the supervisor module running Cisco MDS SAN-OS in the Location or Address field of your browser. You see the installation page displayed by the HTTP server of the supervisor module.

When you connect to the server for the first time, it checks to see if you have the correct Sun Java Virtual Machine version installed on your workstation. If you do not have the correct version installed, a link is provided to the appropriate web page on the Sun Microsystems website so you can install it.

a. Click the **Sun Java Virtual Machine** software link (if required) to install the software.

b. Using the instructions provided by the Sun Microsystems website, reconnect to the supervisor module by reentering the IP address or host name in the Location or Address field of your browser.

Note Fabric Manager requires Java version 1.4(x). We recommend a minimum of Java version 1.4.2. To use IPv6 addresses, you must have a minimum of Java version 1.5. To change the Java Runtime Environment (JRE) version, start **Java Web Start** and set the Java preferences.

Step 2 Click the **Fabric Manager Web Services** installation link. You see a prompt asking for permission to install the application on your workstation.

Step 3 Click **Yes** to run the installer, which detects the installed version of the software, and prompts for upgrades or downgrades and other options if applicable.

Note If TCP port 80 is in use, Fabric Manager Web Services checks port 8080 next. If that port is also in use, Fabric Manager Web Services uses the next available port. You can set the TCP port that you want Fabric Manager Web Services to use during the installation process.

Unless you specify a different directory on a Windows PC, the software is installed in the default location of **C:\Program Files\Cisco Systems\MDS 9000**. A **Cisco MDS** 9000 program group is created under Start > Programs. This program group contains shortcuts to Fabric Manager and Device manager.

On a UNIX (Solaris or Linux) machine, the installation path is /usr/local/cisco_mds9000. If this directory is not writable by the user, which is the case for non-root users, the default is set to $HOME/cisco_mds9000. Shell scripts are created in the bin directory.

Note On a Windows PC, you install Fabric Manager Web Services as a service. This service can then be administered using the Services Panel from the Windows Control Panel. By default, Fabric Manager Web Services automatically starts when the workstation is rebooted. You can change this behavior by modifying the properties in the Services Panel.

Using Fabric Manager Web Services with SSL

Fabric Manager Web Services uses TCP port 80 by default. If you want to install SSL certificates and use Fabric Manager Web Services over HTTPS (using TCP port 443 or another custom port), you need a certificate for each external IP address that accepts secure connections. You can purchase these certificates from a well-known Certificate Authority (CA).

To modify Fabric Manager Web Services to use SSL, follow these steps:

Step 1 Stop Fabric Manager Web Services if you have already launched it. If you installed this on Windows, you can stop the service using Windows Services under Administrative Tools.

Step 2 Use a text editor to open \tomcat\conf\server.xml from the directory where Fabric Manager Web Services is installed. You see the following lines in the beginning after some copyright information:

```
<Connector className="org.apache.catalina.connector.http.HttpConnector"
           port="80" minProcessors="5" maxProcessors="75"
           enableLookups="false" redirectPort="8443"
           acceptCount="10" debug="0" connectionTimeout="60000"/>
<!-- Define an SSL HTTP/1.1 Connector on port 8443 -->
<!--
<Connector className="org.apache.catalina.connector.http.HttpConnector"
           port="8443" minProcessors="5" maxProcessors="75"
           enableLookups="true"
           acceptCount="10" debug="0" scheme="https" secure="true">
  <Factory className="org.apache.catalina.net.SSLServerSocketFactory"
           clientAuth="false" protocol="TLS"/>
</Connector>
-->
```

Step 3 Comment the first <Connector> element and uncomment the second one. Note that the port changes from 8443 to 443 and keystore and keypass are added. Your file should look like the following example:

```
<!--
<Connector className="org.apache.catalina.connector.http.HttpConnector"
           port="80" minProcessors="5" maxProcessors="75"
           enableLookups="false" redirectPort="8443"
           acceptCount="10" debug="0" connectionTimeout="60000"/>
-->
<!-- Define an SSL HTTP/1.1 Connector on port 8443 -->
<Connector className="org.apache.catalina.connector.http.HttpConnector"
           port="443" minProcessors="5" maxProcessors="75"
           enableLookups="true"
           acceptCount="10" debug="0" scheme="https" secure="true">
  <Factory className="org.apache.catalina.net.SSLServerSocketFactory"
           clientAuth="false" protocol="TLS"
           keystoreFile="C:\Program Files\Cisco Systems\MDS 9000\keystore"
           keystorePass="changeit"/>
</Connector>
```

Step 4 Save this file.

Step 5 Restart **Fabric Manager Web Services**.

Launching Fabric Manager Web Services

Before you can use Fabric Manager Web Services to monitor a switch, the service must be started on the server you are connecting through. The browser does not have to be on the same workstation where Fabric Manager Web Services is installed.

To launch Fabric Manager Web Services, follow these steps:

Step 1 If you are on the same workstation where you installed Fabric Manager Web Services, then open your browser and in the Location field enter **http://localhost:*PORT***. Enter your port number if you specified a different port during installation. You can omit the port number if you used port 80 by default.

If you are on a different workstation from where you installed Fabric Manager Web Services, then open your browser and in the Location field enter **http://<*yourServerAddress*>:*PORT***, where <*yourServerAddress*> is the address where you installed Fabric Manager Web Services, and *PORT* is 80 by default. Enter your port number if you specified a different port during installation.

Tip Click the Windows XP **Start > Control Panel > Administrative Tools > Services** to verify that Fabric Manager Web Services is started. To start **Fabric Manager Web Services**, use a browser to go to the location of the service.

You can also view this information using the **Admin > Status** menu of the Fabric Manager Web Services client.

On a UNIX workstation, use the following command:

```
$ /usr/local/cisco_mds9000/bin/FMWebClient.sh status
```

You see the login screen for Fabric Manager Web Services (see Figure 6-2). The text field at the bottom shows the Message of the Day from the server you logged into.

Figure 6-2 Fabric Manager Web Services Login Screen

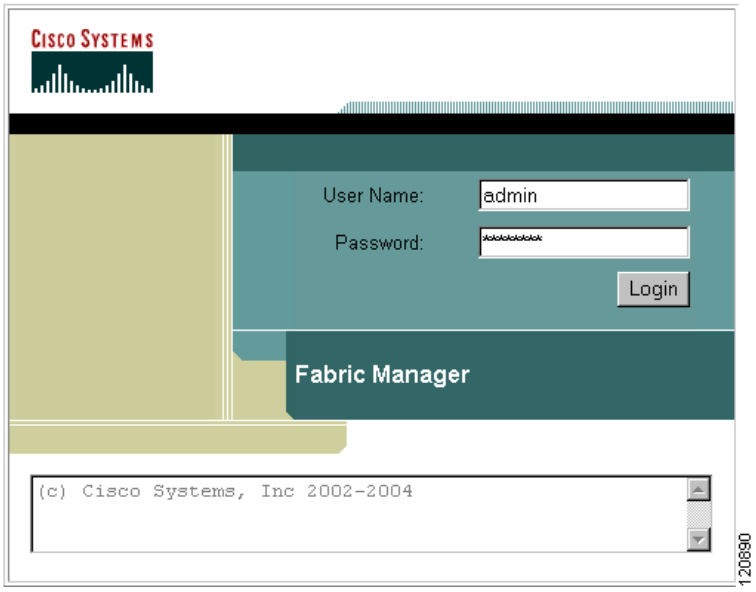

Step 2 Enter your user name and password.

Step 3 Click **Login**.

Note When attempting to log in, you may see **No valid ID - server cannot find ID**. This happens if no discovery is done for that fabric using Fabric Manager Server, so no user name or password is recorded. To resolve this issue, open Fabric Manager and discover the fabric. Performance Manager will record the user name and password. Then log in to Fabric Manager Web Services again.

After launching Fabric Manager Web Services, you see the initial screen, which you can also see by clicking the **Events** tab followed by the **Summary** tab (see Figure 6-3). Fabric Manager Web Services polls the Fabric Manager Server database to display the managed devices in the left pane.

Figure 6-3 Events > Summary Screen

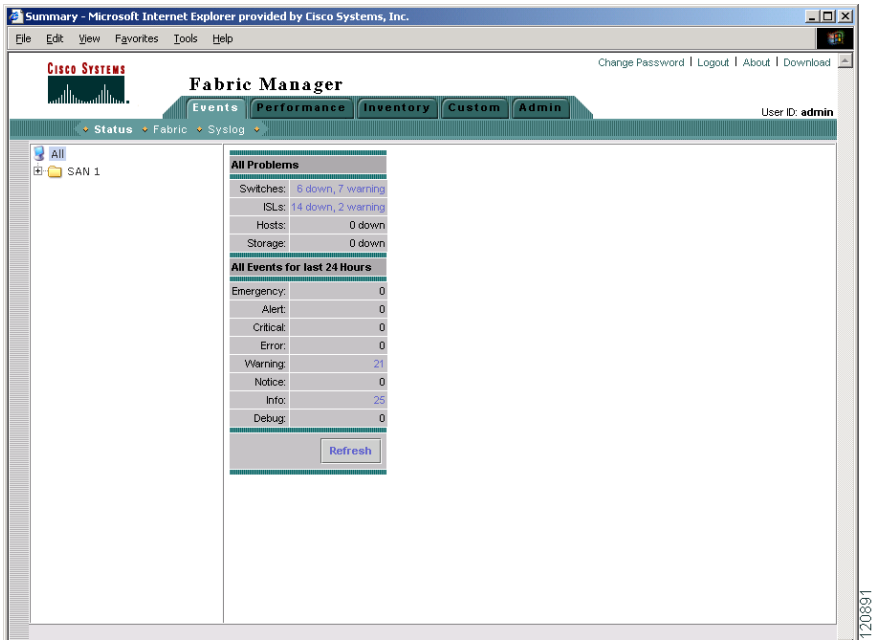

Navigating Fabric Manager Web Services

With most screens, Fabric Manager Web Services has standardized on certain navigation conventions.

Navigation Tree

You can use the filter navigation tree on the left pane to access the areas you want. In general:

- Select **SAN** to view information for all fabrics and VSANs in the SAN. When you do this, a "Fabric" column is added as the first column of the tables.
- Click a fabric folder to view information for that specific fabric.
- Some screens have expandable fabric folders. You can expand the fabric folders (by clicking the + or - icons in front of the folders) to see a list of VSANs in that fabric. Select a VSAN to view information for that VSAN.

The features accessible from the tabs are limited to the areas you select in the filter tree.

Table Filtering and Navigation

You can filter the display of some tables to view subsets of the information. At the top right of these tables are one or more drop-down lists. Select an item from these lists, and then click **Filter** to filter the table information on that item.

You can change the number of rows displayed per page by selecting a number from the Rows per page drop-down list at the lower left corner of the table. Once you select a number, the table is updated with the new number of rows–you do not have to click a button.

For tables with multiple pages of information, you can:

- Jump to the first or last page of the table by clicking the “first page” or “last page” icons (arrows with a bar in front of it)
- Jump to the next page or previous page by clicking the “next page” or “previous page” icons (arrows)
- Jump to a specific page by entering the page number in the Go to page field and clicking the **Go** button.

You can search certain columns in the tables for information if a table column has a black icon next to the column head. Click the icon to display a Search dialog box.

Printing

There is a **Print** icon in the lower right corner of some tables. Click this icon to view the table in a printer-friendly format. You can then print the page from the browser.

Exporting to a File

There is an **Export** icon in the lower right corner of some tables. Click this icon to export the data to a .CSV file that can be read by programs such as Microsoft Excel.

Admin Tab

Only users with the **network-admin** role can see the Admin tab. Users configured with **network-operator** or an empty role cannot see the Admin tab.

Recovering a Web Services Password

Fabric Manager Web Services user passwords are encrypted and stored locally on the workstation where you installed Web Services. If you forget a password, you can make a new network-admin user locally on the workstation where you installed Web Services and then log in and delete the old user account under the Admin tab.

To create a user on the workstation where you installed Web Services and delete the old user, follow these steps:

Step 1 Go to the Web Services installation directory and **cd** to the bin directory.

Step 2 Enter the following line to create a user:

```
webAddUser <userName> <password>
```

Step 3 Stop Fabric Manager Web Services if it is running. If you installed this on Windows, you can stop the service using Windows Services under Administrative Tools.

Step 4 Launch **Fabric Manager Web Services**.

Step 5 Click **Admin > Configure > Web Users > Local Database**.

You see the list of users in the local database.

Step 6 Select the user that you want to delete and click **Delete** to remove the old user.

Events

The Events tab shows events and issues for the selected items, persistent across user sessions.

The Events tab contains the following subtabs:

- Summary—Shows a summary of events and problems for all SANs, or a selected SAN, fabric, or switch. You can click any of the blue links for more information about that item.
- Fabric—Shows a detailed list of events and hardware, or accounting. You can filter these events by severity, date, and type of event.
- Syslog—Shows a detailed list of system messages. You can filter these events by severity, date, and type of event.

Viewing Summary Information

To view a summary of events and problems using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Events** tab followed by the **Summary** tab.

You see the Summary screen. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server. In the right pane is a summary table of problems and events for the last 24 hours shown in Figure 6-4.

Figure 6-4 Events Summary Screen

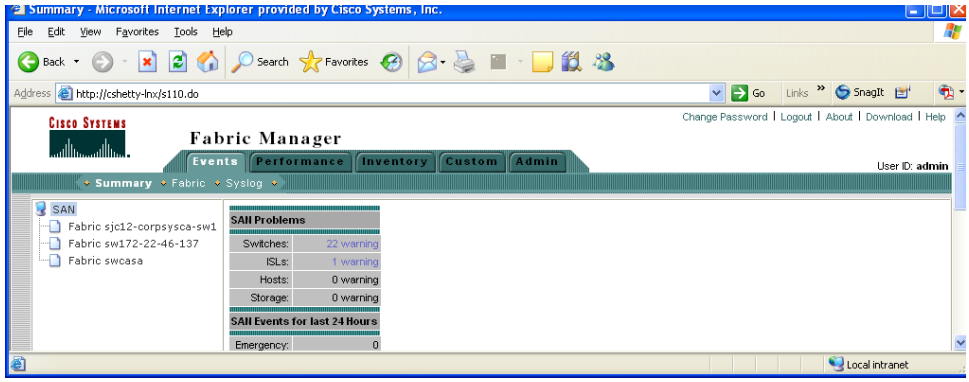

Step 2 Do one of the following:

- Select **SAN** to display summary information for all fabrics.
- Select one of the fabrics to display summary information for that fabric.

Step 3 Click the warnings next to Switches, ISLs, Hosts, or Storage (other than 0) to see an inventory of switches, ISLs, or end devices for that fabric.

Step 4 Click the number of events next to the event severity levels (Emergency, Alert, Critical, Error, Warning, Notice, Info, or Debug) to see a table of events and descriptions for that fabric.

Viewing Fabric Information

To view a detailed list of events and hardware, or accounting using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Events** tab followed by the **Fabric** tab.

You see the Fabric screen shown in Figure 6-5. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server.

Figure 6-5 Fabric Events for Entire SAN

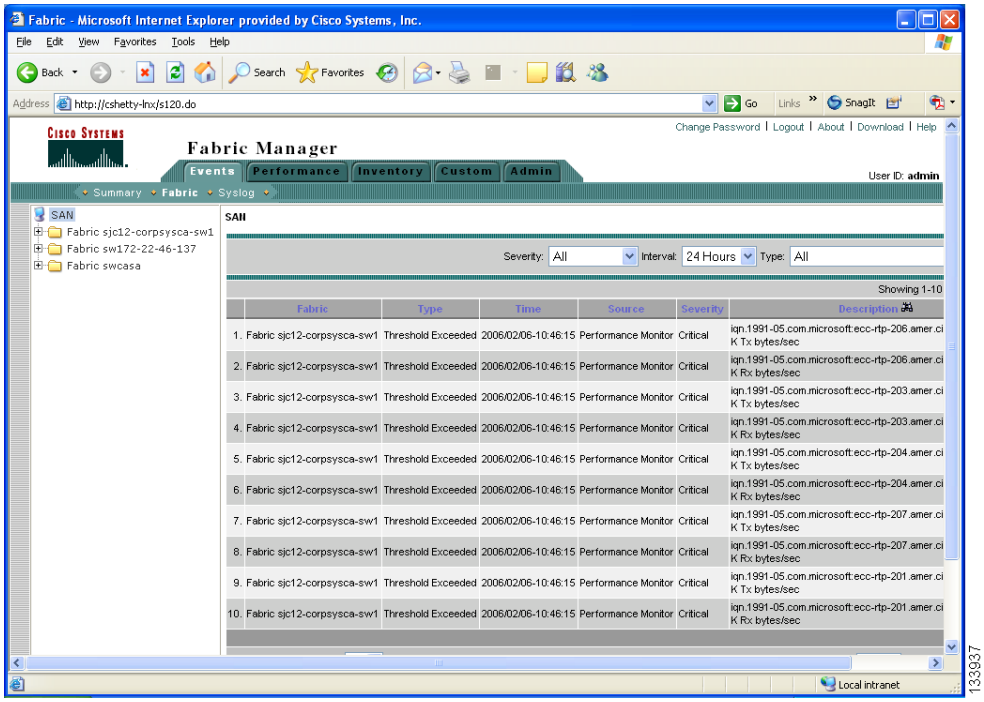

Step 2 Do one of the following:

- Click **SAN** to display event information for all fabrics (see Figure 6-5).
- Click one of the fabrics to display event information for that fabric.
- Expand a fabric and click one of the switches to display event information for that switch.

Note Click a column head to sort the events by that column.

Viewing Syslog Information

To view a detailed list of system messages using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Events** tab followed by the **Syslog** tab.

You see the Syslog screen. In the left navigation pane you see a list of fabrics monitored by Fabric Manager Server shown in Figure 6-6.

Figure 6-6 Syslog Events for Entire SAN

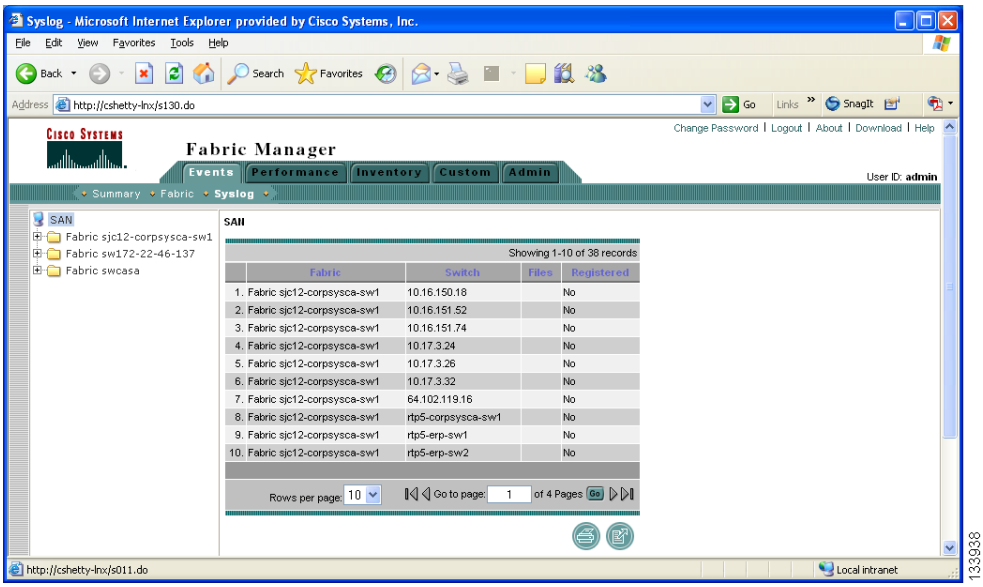

Step 2 Do one of the following:

- Select **SAN** to display a table of syslog information for all fabrics (see Figure 6-6).
- Select one of the fabrics to display a table of syslog information for that fabric.
- Expand a fabric and select one of the switches to display syslog information for that switch.

Step 3 If you have selected a fabric and one or more switches in that fabric have system messages, you see Events, Hardware, Accounting, and Link Incidents in the Files column. Click one of these message types to see system messages for the switches in that fabric filtered by the message type you clicked.

Note Click a column head to sort the system messages by that column.

If you have selected a switch, choose an interval and a message type from the drop-down lists, and then click **Filter** to see system messages filtered by the message type you chose.

Performance

The Performance tab shows an overview of the average throughput and link utilization of SAN components. You see pie charts for the throughput and utilization. You can click a pie chart to view a table of the data. In these tables, clicking a blue link displays a graph of that data, if applicable. The Filter drop-down list at the top right of the screen allows you to filter the data based on various periods of time.

The Performance tab contains the following subtabs:

- Summary—Shows the total utilization and throughput in summary form.
- End Devices—Shows a detailed list of end devices (host or storage), port traffic, and errors.
- ISLs—Shows a detailed list of ISL traffic and errors.
- Flows—Shows a detailed list of host-to-storage traffic.
- Traffic Analyzer—Shows a summary of SPAN ports configured in the SAN and any traffic analyzers configured.
- Prediction—Displays a graph that predicts future performance to help determine when storage network connections will become overutilized.

Note Performance Manager shows statistics for fabrics that you have configured collections for using the Collection Wizard.

Viewing Performance Summary Information

To view total utilization and throughput in summary form using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Performance** tab followed by the **Summary** tab.

You see the Summary screen. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server (see Figure 6-7).

Figure 6-7 Performance Summary for Entire SAN

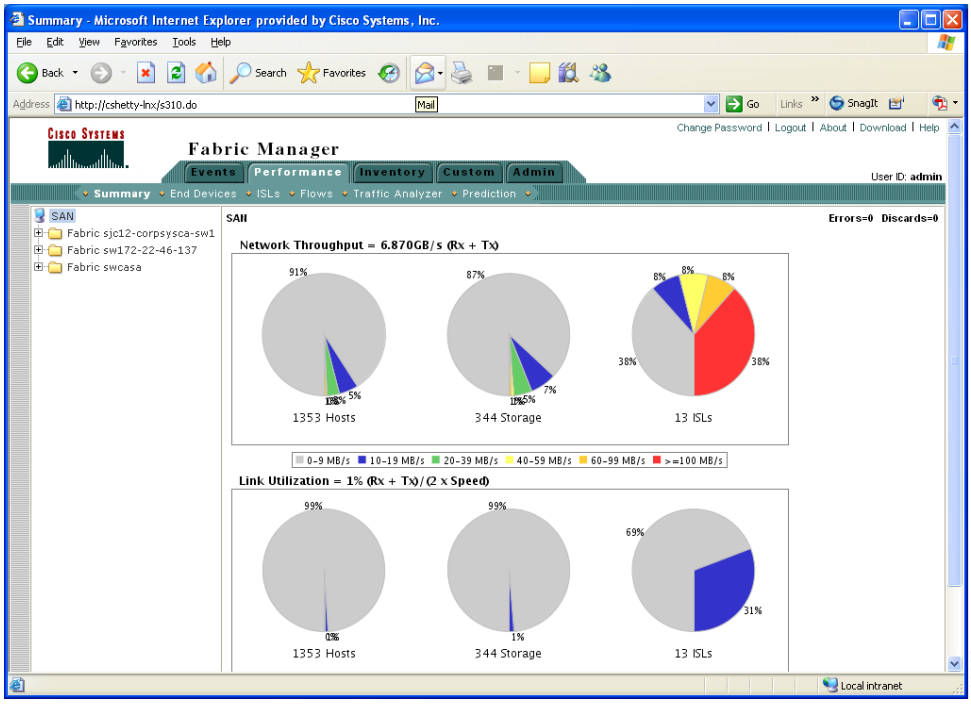

Step 2 Do one of the following:

- Select **SAN** to display network throughput and link utilization information for all fabrics.
- Select one of the fabrics to display network throughput and link utilization information for that fabric.
- Expand a fabric and select one of the VSANs to display network throughput and link utilization information for that VSAN.

Note Click a pie chart (Hosts, Storage, or ISLs) to go to the appropriate performance table.

Viewing Performance Information for End Devices

To view host and storage port traffic and errors using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Performance tab** followed by the **End Devices** tab.

You see the End Devices screen shown in Figure 6-8. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server.

Figure 6-8 End Devices for Entire SAN

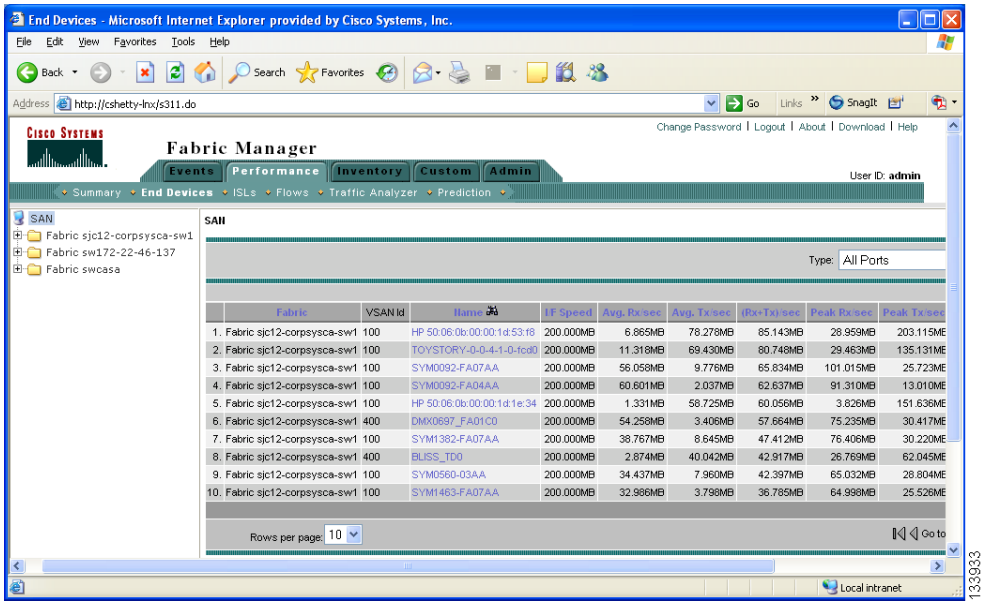

Step 2 Do one of the following:

- Select **SAN** to display performance information for the end devices in all fabrics in the SAN (see Figure 6-8).
- Select one of the fabrics to display performance information for the end devices for all VSANs in that fabric.
- Expand a fabric and select one of the VSANs to display performance information for the end devices in that VSAN.

Note There are variations to this procedure. In addition to these basic steps, you can also:

- Click a column head to sort the performance information by that column.
- Select the type of ports and the time range, and click **Filter** to filter the display of those criteria.
- Click the name of a device in the Name column to see a graph of the traffic on that device for the past 24 hours. You can change the time range for this graph by selecting it from the drop-down list in the upper right corner.

- If you select the Host Enclosure type filter, you see an icon in front of the links in the Name column. Click the icon to display a new table showing the detailed information about that host.

Viewing Performance Information for ISLs

To view ISL traffic and errors using Fabric Manager Web Services, follow these steps:

Step 1 Select the **Performance** tab followed by the **ISLs** tab.

You see the ISLs screen. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server.

Figure 6-9 ISL Performance for Entire SAN

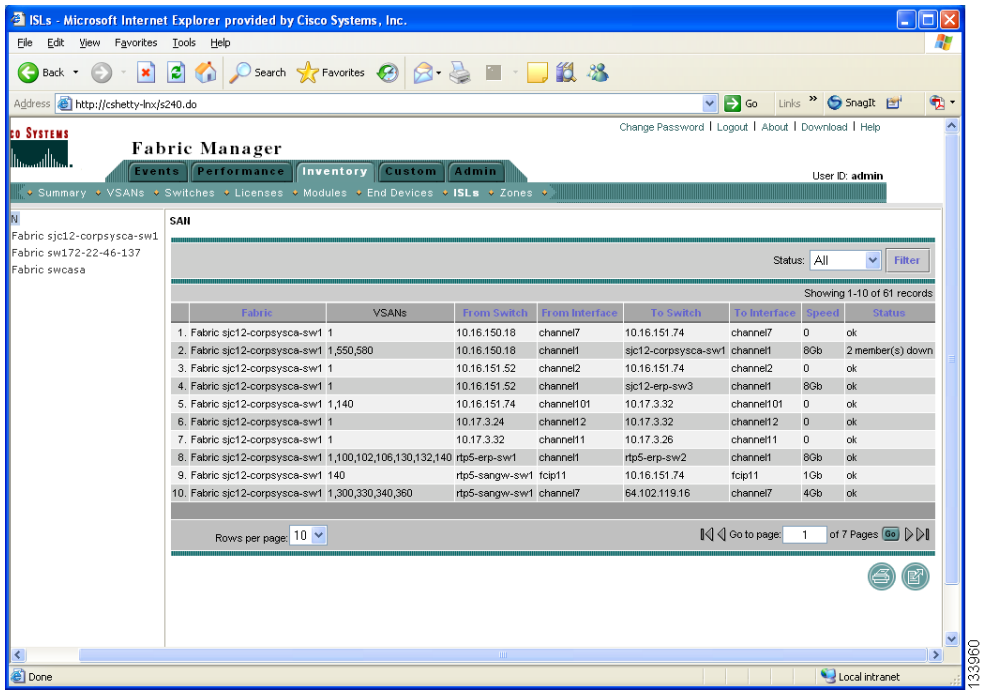

Step 2 Do one of the following:

- Select **SAN** to display performance information for the ISLs in all fabrics in the SAN (see Figure 6-9).
- Select one of the fabrics to display performance information for the ISLs for all VSANs in that fabric.
- Expand a fabric and select one of the VSANs to display performance information for the ISLs in that VSAN.

Note There are variations to this procedure. In addition to the basic steps described above, you can also:

- Click a column head to sort the performance information by that column.
- Select the time range, and click **Filter** to filter the display.
- Click the name of an ISL in the Name column to see a graph of the traffic across that ISL for the past 24 hours. You can change the time range for this graph by selecting it from the drop-down list in the upper right corner.

Viewing Performance Information for Flows

To view host and storage traffic using Fabric Manager Web Services, follow these steps:

Step 1 Choose the **Performance** tab followed by the **Flows** tab.

You see the Flows screen shown in Figure 6-10. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server.

Figure 6-10 Performance Flows for the Entire SAN

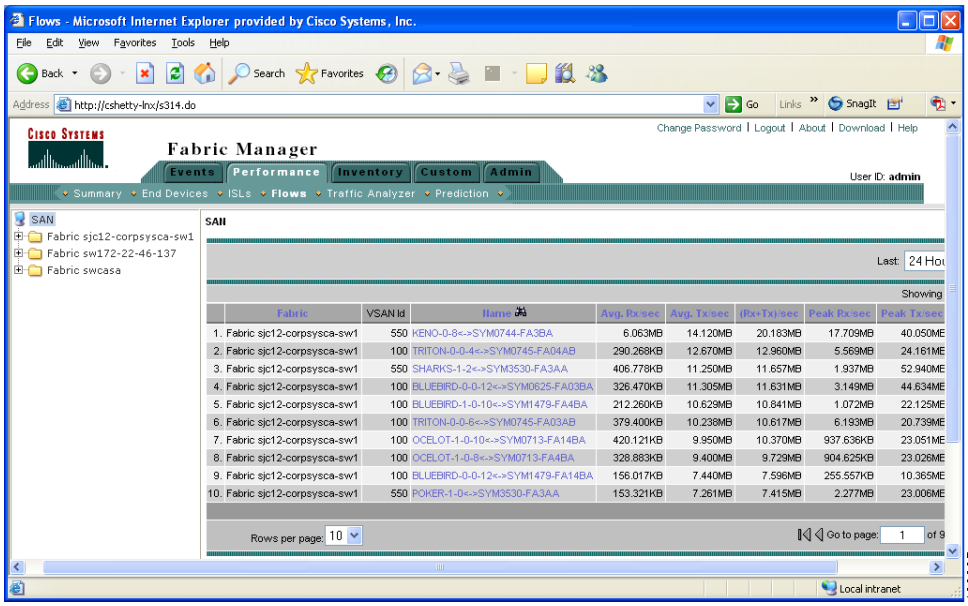

Step 2 Do one of the following:

- Select **SAN** to display performance information for the flows in all fabrics in the SAN (see Figure 6-10).
- Select one of the fabrics to display performance information for the flows for all VSANs in that fabric.
- Expand a fabric and select one of the VSANs to display performance information for the flows in that VSAN.

Note There are variations to this procedure. In addition to these basic steps, you can also:

- Click a column head to sort the performance information by that column.
- Select the time range, and click **Filter** to filter the display.
- Click the name of a flow in the Name column to see a graph of the traffic over that flow for the past 24 hours. You can change the time range for this graph by selecting it from the drop-down list in the upper right corner.

Viewing Detailed Traffic Information

To view SPAN port detailed traffic using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Performance** tab followed by the **Traffic Analyzer** tab.

You see the Traffic Analyzer screen shown in Figure 6-11. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server.

Figure 6-11 Performance Traffic for Entire SAN

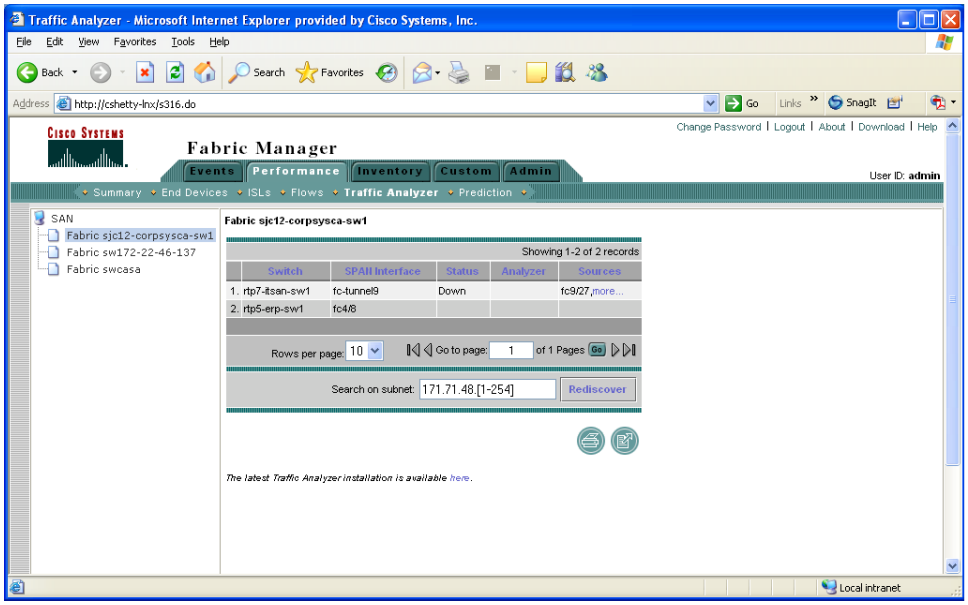

Step 2 Do one of the following:

- Select **SAN** to display a list of SPAN ports for switches in all fabrics in the SAN (see Figure 6-11).
- Select one of the fabrics to display a list of SPAN ports for switches in that fabric.

Note Click a column head to sort the traffic analyzer information by that column.

Viewing Predicted Future Performance

To plan storage network changes, it is necessary to determine when configuration changes (such as re-zoning) may be needed to meet growing performance demands. Fabric Manager Server provides a performance prediction report to enable you to more easily predict when storage network connections will become overutilized.

In general, to create a performance prediction report, do the following:

- Specify the period of time in the past that you want to use as a sample to predict the future performance.
- Specify the threshold values that you do not want to exceed.
- Specify the period of time in the future for which you want to view performance.

Fabric Manager Server extrapolates the performance and lists in chronological order which interfaces are expected to reach the threshold within the specified time period.

Using the Default Values

When you first view predicted future performance by selecting **Performance > Prediction**, you see a table showing the predicted performance for your entire SAN using the default values. The default values are:

- Scope—entire SAN
- Past performance period—Month
- Future performance period—Month
- Threshold—80%
- SAN elements or links—ISLs
- Performance prediction type—Average

Click a link in the Name column to view a graph of that ISL's performance for the past 24 hours. To view the performance for the past week, month, year, or custom time, select an option from the drop-down list.

Using Your Own Values

To view a table of predicted future performance with your own values using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Performance** tab followed by the **Prediction** tab.

You see the Prediction screen shown in Figure 6-12. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server. In the right pane you see a table showing the predicted performance for the SAN using the default values.

Figure 6-12 Performance Prediction for an Entire SAN

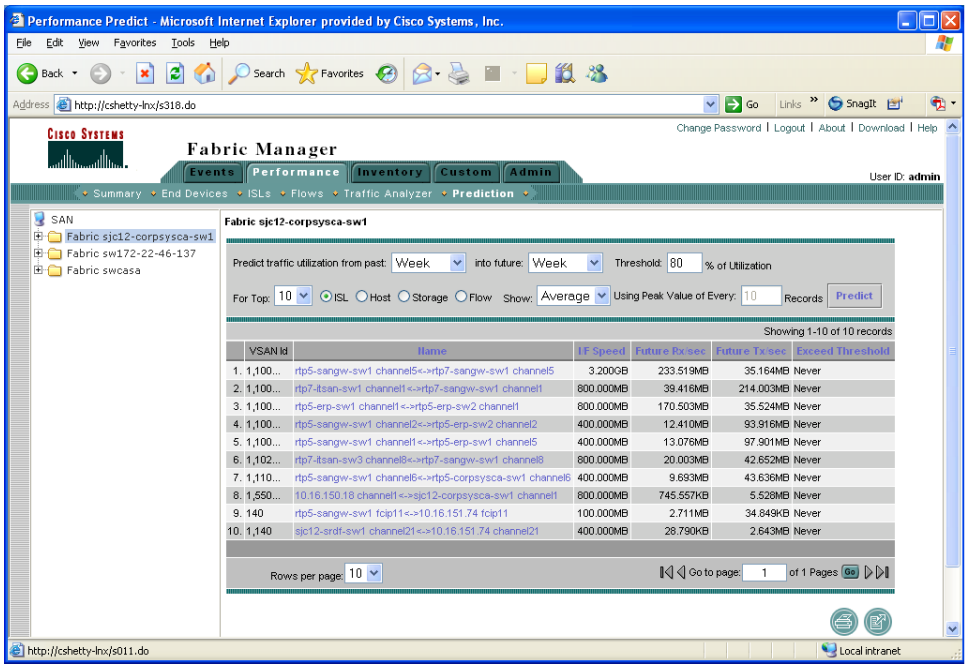

Step 2 Do one of the following:

- Select **SAN** to specify that the prediction report will be generated for all fabrics in the SAN (see Figure 6-12).
- Select one of the fabrics to specify that the prediction report will be generated for all VSANs in that fabric.
- Expand a fabric and select one of the VSANs to specify that the prediction report will be generated for that VSAN.

Step 3 Select the period of time (**Week**, **Month**, **3 Months**, **6 Months**, and **Year**) to use to predict performance from the "past" drop-down list.

Step 4 Select the period of time (**Week**, **Month**, **3 Months**, **6 Months**, and **Year**) for which to make the prediction from the "future" drop-down list.

Step 5 Enter the threshold percentage (1-100) of utilization that you do not want the traffic to exceed.

Step 6 Enter the number of ISLs, hosts, storage devices, or flows for which you want to make the prediction. The prediction will show the top **10**, top **20**, or top **50** with the most traffic.

Step 7 Select the type of traffic prediction to show:

- **Average**—The average value of all the sample data is used.
- **Peak**—The average value of all the peak values is used. The number of peak values is obtained by dividing the total number of records into groups based on the number you enter in the Use Peak Value of Every *xx* Records field. For example, if you have 1000 records and you enter 100 into the field, your records are divided into 10 groups and 10 peak values are used.

Step 8 Click **Predict.**

You see the prediction table with the new data. Click the links in the Name column to show performance charts based on the history data.

Note You can also click a column head to sort the performance information by that column.

Inventory

The Inventory tab shows an inventory of the selected SAN, fabric, or switch. You can export this information to an ASCII file in comma-separated value format that can be read by applications such as Microsoft Excel. You can set the number of rows and columns per page.

The Inventory tab contains the following subtabs:

- Summary—Shows VSANs, switches, ISLs, ports, and end devices.
- VSANs—Shows details about VSANs.
- Switches—Shows details about switches.
- Licenses—Shows details about the licenses in use in the fabric.
- Modules—Shows details for MDS switching and services modules, fans, and power supplies.
- End Devices—Shows the host and storage ports.
- ISLs—Shows the Inter-Switch Links.
- Zones—Shows the active zone members (including those in inter-IVR VSAN zones).

Viewing Summary Inventory Information

To view a summary of VSANs, switches, ISLs, ports, and end devices using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Inventory** tab followed by the **Summary** tab.

You see the Summary screen shown in Figure 6-13. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server.

Figure 6-13 Inventory Summary for all Fabrics

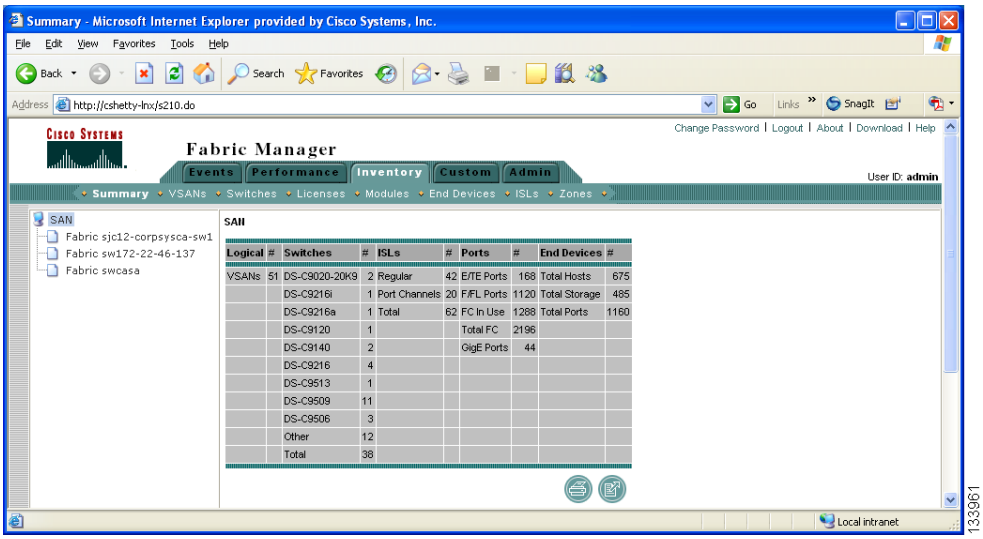

Step 2 Do one of the following:

- Select **SAN** to display a summary of inventory information for all fabrics in the SAN (see Figure 6-13).
- Select one of the fabrics to display a summary of inventory information for that fabric.

Viewing Detailed Information for VSANs

To view detailed inventory information about VSANs using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Inventory** tab followed by the **VSANs** tab.

You see the VSAN inventory screen shown in Figure 6-14. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server.

Figure 6-14 VSAN Inventory Information for all Fabrics in the SAN

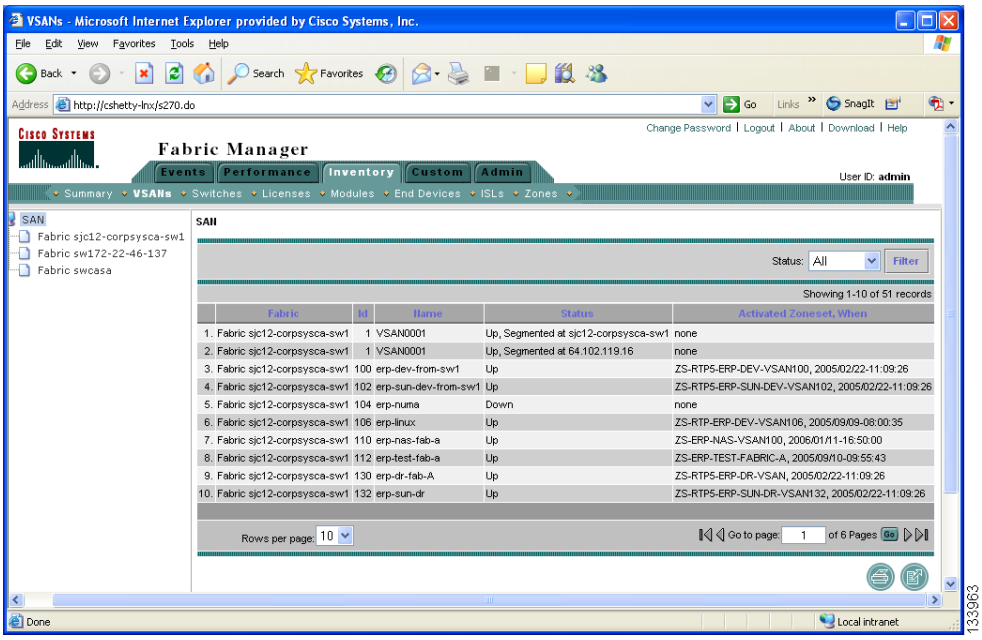

Step 2 Do one of the following:

- Select **SAN** to display VSAN inventory information for all fabrics in the SAN (see Figure 6-14).
- Select one of the fabrics to display VSAN inventory information for that fabric.

Note There are variations to this procedure. In addition to these basic steps, you can also:

- Click a column head to sort the inventory information by that column.
- Select the status level, then click **Filter** to filter the display to show all VSANs or just those with errors.

Viewing Detailed Information for Switches

To view detailed inventory information about switches using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Inventory** tab followed by the **Switches** tab.

You see the Switches inventory screen shown in Figure 6-15. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server.

Figure 6-15 Switch Inventory for all Fabrics in the SAN

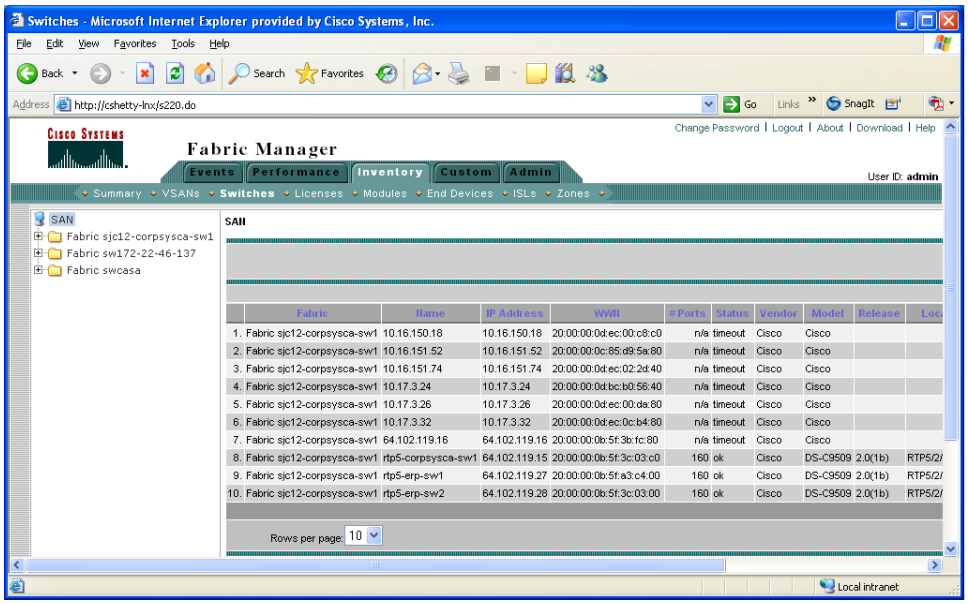

Step 2 Do one of the following:

- Select **SAN** to display switch inventory information for all fabrics in the SAN (see Figure 6-15).
- Select one of the fabrics to display switch inventory information for that fabric.
- Expand a fabric and select one of the VSANs to display switch inventory information for that VSAN.

Note There are variations to this procedure. In addition to these basic steps, you can also:

- Click a column head to sort the inventory information by that column.
- Select the status level, then click **Filter** to filter the display to show all switches or just those with errors.

Viewing License Information

To view license information for switches using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Inventory** tab followed by the **Licenses** tab.

You see the Licenses inventory screen shown in Figure 6-16. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server.

Figure 6-16 License Inventory for Switches in all Fabrics in the SAN

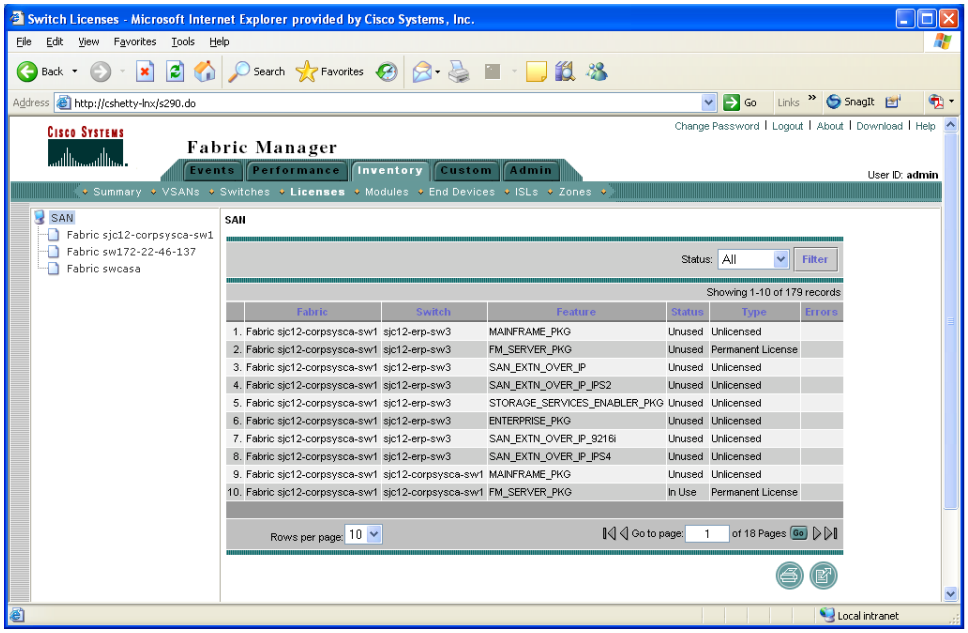

Step 2 Do one of the following:

- Select **SAN** to display license information for switches in all fabrics in the SAN (see Figure 6-16).
- Select one of the fabrics to display license information for switches in that fabric.

Note There are variations to this procedure. In addition to these basic steps, you can also:

- Click a column head to sort the inventory information by that column.
- Select the status level, and click **Filter** to filter the display to show all licenses or just those with errors.

Viewing Detailed Information for Modules

To view detailed inventory information about modules using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Inventory** tab followed by the **Modules** tab.

You see the Modules inventory screen shown in Figure 6-17. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server.

Figure 6-17 Module Inventory for Switches in all Fabrics in the SAN

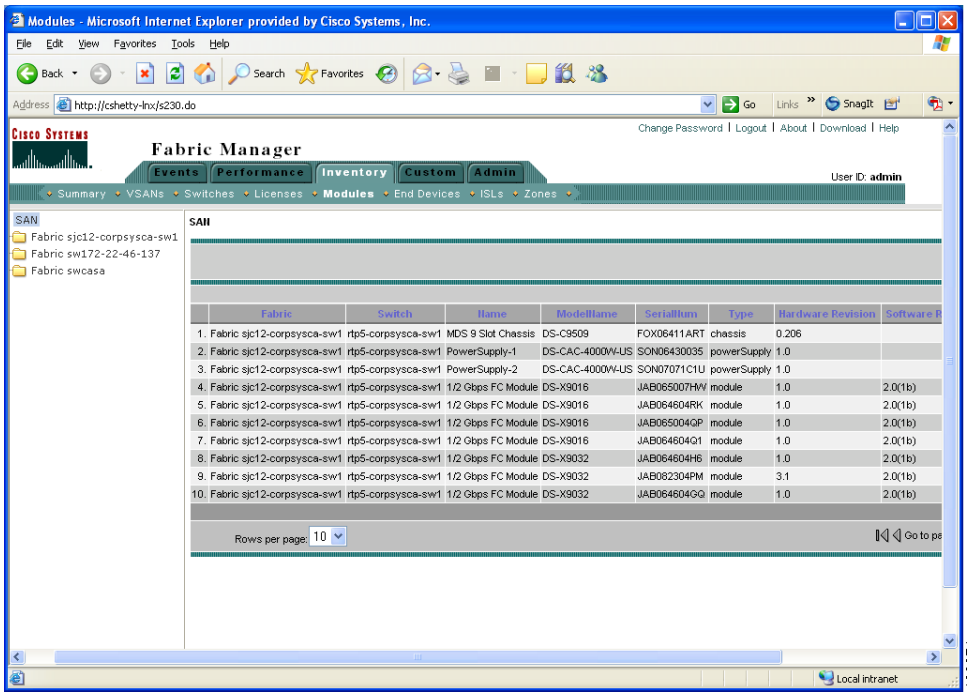

Step 2 Do one of the following:

- Select **SAN** to display module inventory information for switches in all fabrics in the SAN (see Figure 6-17).
- Select one of the fabrics to display module inventory information for switches in that fabric.
- Expand a fabric and select one of the VSANs to display module inventory information for switches in that VSAN.

Note There are variations to this procedure. In addition to these basic steps, you can also:

- Click a column head to sort the inventory information by that column.
- Select the status level, and click **Filter** to filter the display to show all modules or just those with errors.

Viewing Detailed Information for End Devices

To view detailed inventory information about end devices using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Inventory** tab followed by the **End Devices** tab.

You see the End Devices inventory screen shown in Figure 6-18. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server.

Figure 6-18 End Device Inventory for all Fabrics in the SAN

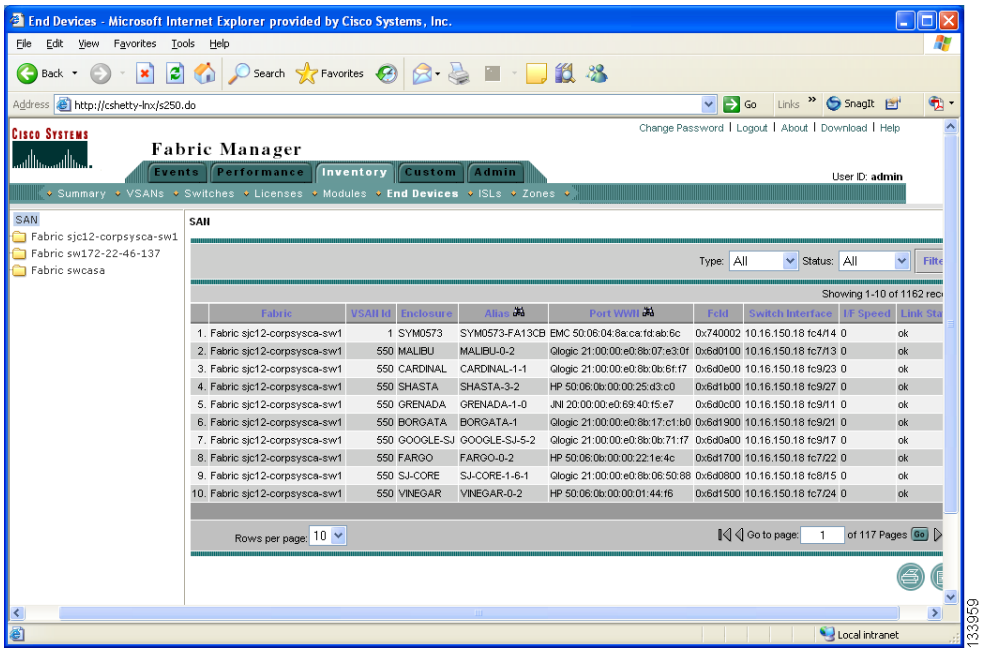

Step 2 Do one of the following:

- Select **SAN** to display end device inventory information for all fabrics in the SAN (see Figure 6-18).
- Select one of the fabrics to display end device inventory information for that fabric.
- Expand a fabric and select one of the VSANs to display end device inventory information for that VSAN.

Note If you filter by hosts or enclosures, you can click a host in the resulting table to see host enclosure performance, a list of hosts, a list of hosts to which your device is connected, and the connection paths. This allows you to see performance statistics for hosts and enclosures.

You can also filter by end devices or by port groups to view aggregate information for those port groups, such as peak and average usage.

Note There are variations to this procedure. In addition to these basic steps, you can also:

- Click a column head to sort the inventory information by that column.

- Select the type and status level, and click **Filter** to filter the display to show all end devices, just hosts, just storage, or just those with errors.

Viewing Detailed Information for ISLs

To view detailed inventory information about ISLs using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Inventory** tab followed by the **ISLs** tab.

You see the ISLs inventory screen shown in Figure 6-19. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server.

Figure 6-19 ISL Inventory for all Fabrics in the SAN

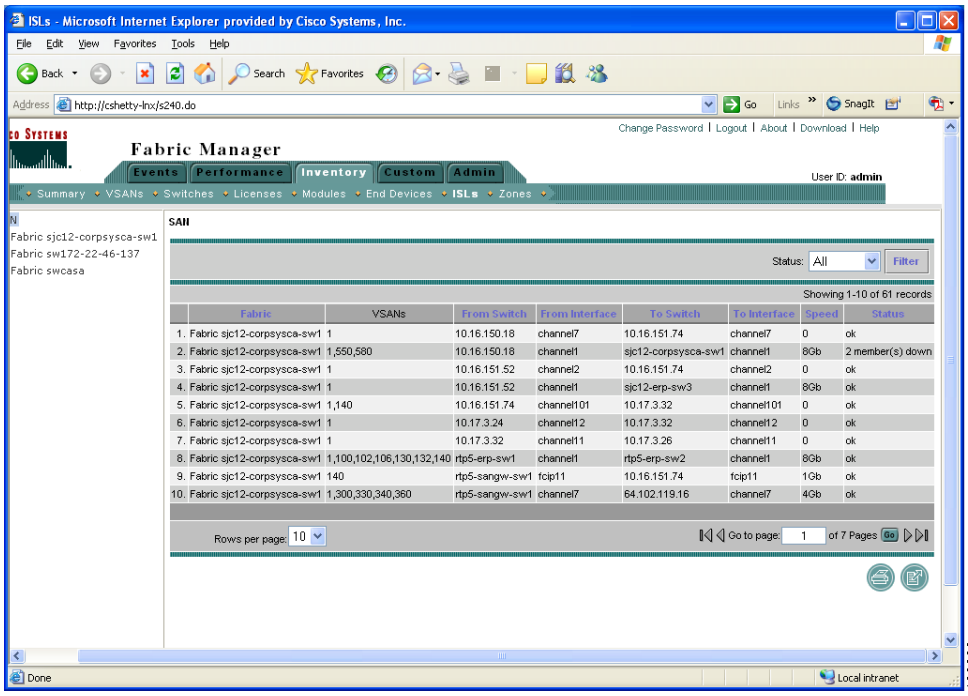

Step 2 Do one of the following:

- Select SAN to display ISL information for all fabrics in the SAN (see Figure 6-19).
- Select one of the fabrics to display ISL inventory information for that fabric.
- Expand a fabric and select one of the VSANs to display ISL inventory information for that VSAN.

Note There are variations to this procedure. In addition to these basic steps, you can also:

- Click a column head to sort the inventory information by that column.
- Select the status level, and click **Filter** to filter the display to show all ISLs or just those with errors.

Viewing Detailed Information for Zones

To view detailed inventory information about zones using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Inventory** tab followed by the **Zones** tab.

You see the Zones inventory screen shown in Figure 6-20. In the left navigation pane you see a list of the fabrics monitored by Fabric Manager Server.

Figure 6-20 Zone Inventory for all Fabrics in the SAN

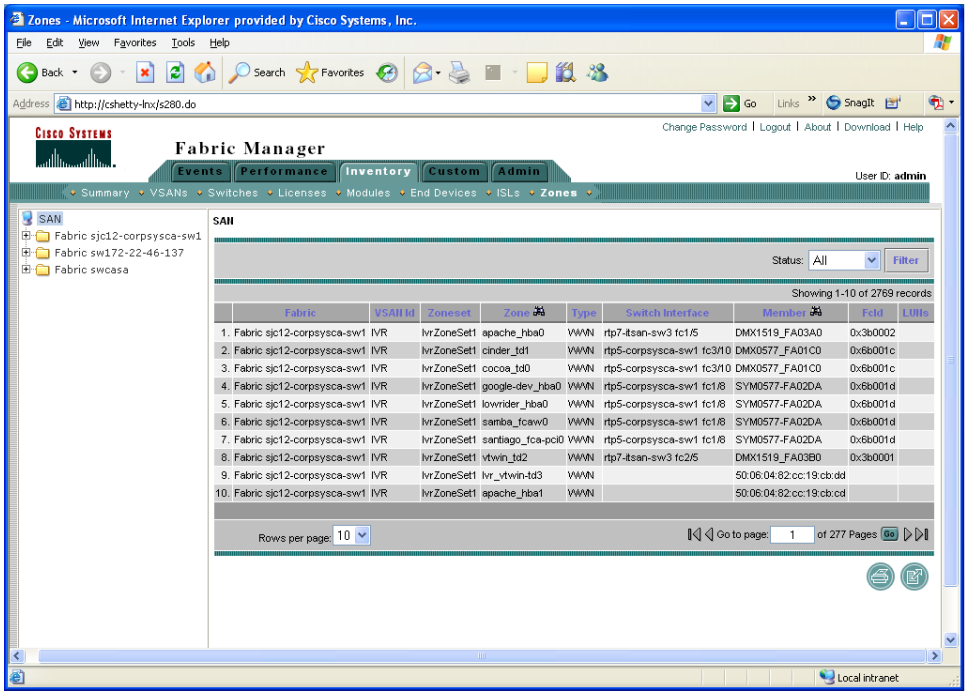

Step 2 Do one of the following:

- Select **SAN** to display zone inventory information for all fabrics in the SAN (see Figure 6-20).
- Select one of the fabrics to display zone inventory information for that fabric.
- Expand a fabric and select one of the VSANs to display zone inventory information for that VSAN.

Note There are variations to this procedure. In addition to these basic steps, you can also:

- Click a column head to sort the inventory information by that column.
- Select the status level, and click **Filter** to filter the display to show all zones or just those with errors.

Custom

The Custom tab allows you to create customized reports based on the historical performance, events, and inventory information gathered by Fabric Manager Server. You can create aggregate reports with summary and detailed views. You can also view previously saved reports.

The Custom tab contains the following subtabs:

- View—Views previously saved reports.
- Generate—Generates a custom report based on the selected report template.
- Edit—Edits an existing report template.
- Create—Creates a report template, allowing you to select any combination of events, performance categories, and inventory.

See the "Creating a Custom Report Template" section on page 6-33.

Listing Custom Reports by Template

Reports you generate are saved by Fabric Manager Server and viewable from the Custom > View tab.

To view a custom report using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Custom** tab followed by the **View** tab.

You see the View Report table shown in Figure 6-21, showing all reports generated and the time you generated the report. You can also navigate to a report in the navigation pane by selecting the report template you used and clicking the report.

Figure 6-21 View Report Table

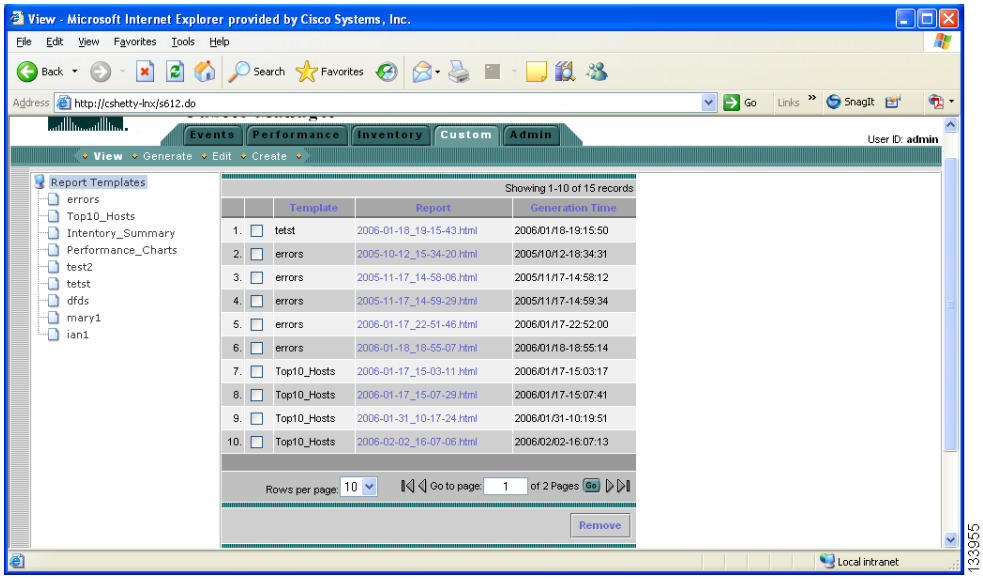

Step 2 Select the report that you want to view from the Navigation pane (see Figure 6-21). You see the report in the main screen.

You see the report in a new browser window if you click the report in the report table.

Generating Custom Reports from a Template

You can generate custom reports from any previously saved report template.

To create a custom report using Fabric Manager Web Services, follow these steps:

Step 1 Select a SAN, fabric, or VSAN on which to base the report.

Step 2 Click the **Custom** tab followed by the **Generate** tab.

You see the Generate Report dialog box shown in Figure 6-22.

Figure 6-22 Generate Custom Report Dialog Box

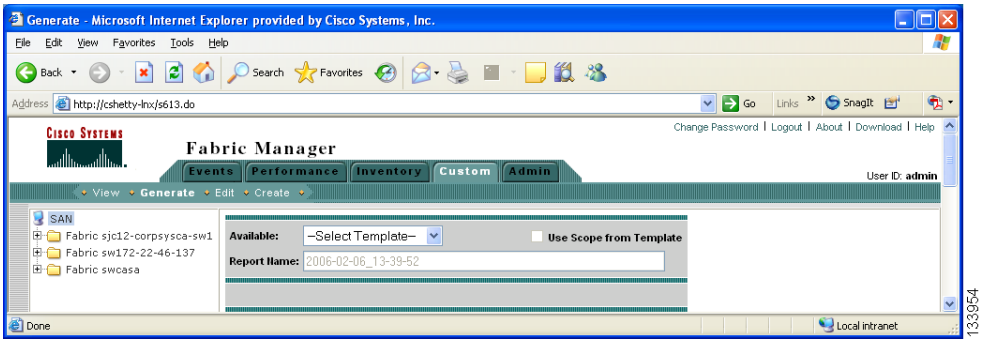

Step 3 Choose a report template from the Available drop-down list.

Step 4 Optionally, change the name of the report. By default, reports are named based on the date and time generated.

Step 5 Optionally, deselect the **Use Scope from Template** check box to override the scope defined by the filter type.

Step 6 Click **Generate** to generate a report based on this template.

After a moment, you see the report results in a new browser window. You can also see the report by clicking **Custom > View** and selecting the report name from the report template you used in the navigation pane.

Creating a Custom Report Template

You can create custom reports from all or any subset of information gathered by Fabric Manager Server. You create a report template by selecting events, performance, and inventory statistics that you want in your report and set the desired SAN, fabric or VSAN to limit the scope of the template. You can generate a report of your fabric based on this template immediately or at a later time. Fabric Manager Web Services saves each report based on the report template used and the time you generate the report. You can view the generated report by clicking **Custom > View** and navigating to the report template and report name.

To create a custom report template using Fabric Manager Web Services, follow these steps:

Step 1 Click **Custom > Create**.

You see the Create Report dialog box shown in Figure 6-23.

Figure 6-23 Create Report Dialog Box

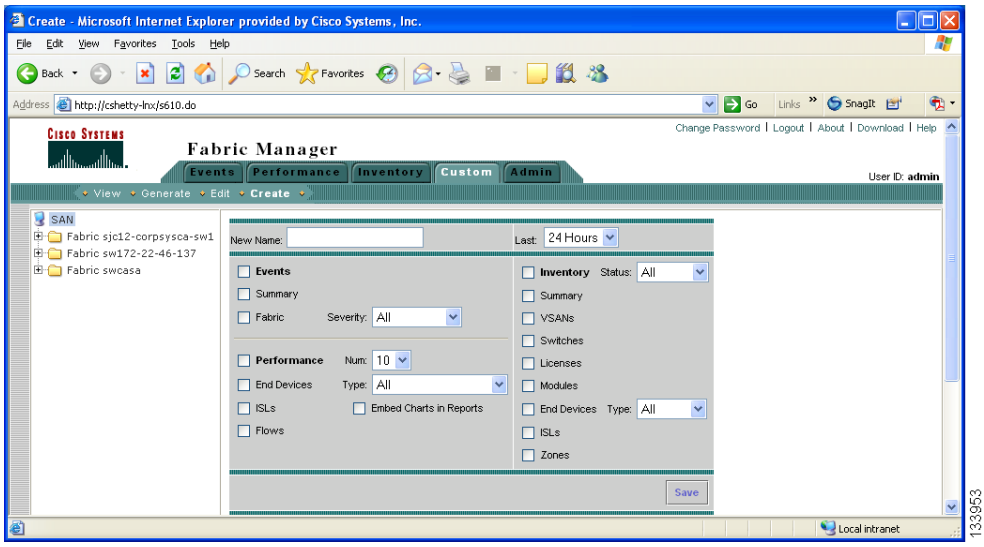

Step 2 Provide a New Name for the report.

Step 3 Indicate the information you want in the report by checking the **Events**, **Performance**, and **Inventory** check boxes.

Step 4 Optionally, select Severity for events, Status for inventory information, or Type of end devices for performance information and inventory information.

Step 5 Click **Save** to save this report template.

Modifying a Custom Report Template

To edit a custom report template using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Custom** tab followed by the **Edit** tab.

You see the Edit Report dialog box

Step 2 Choose a report template from the Available drop- down list and click **Open**.

You see the current information that this report gathers. The Top Ten Report is shown in Figure 6-24.

Figure 6-24 Information Gathered by the Top Ten Report

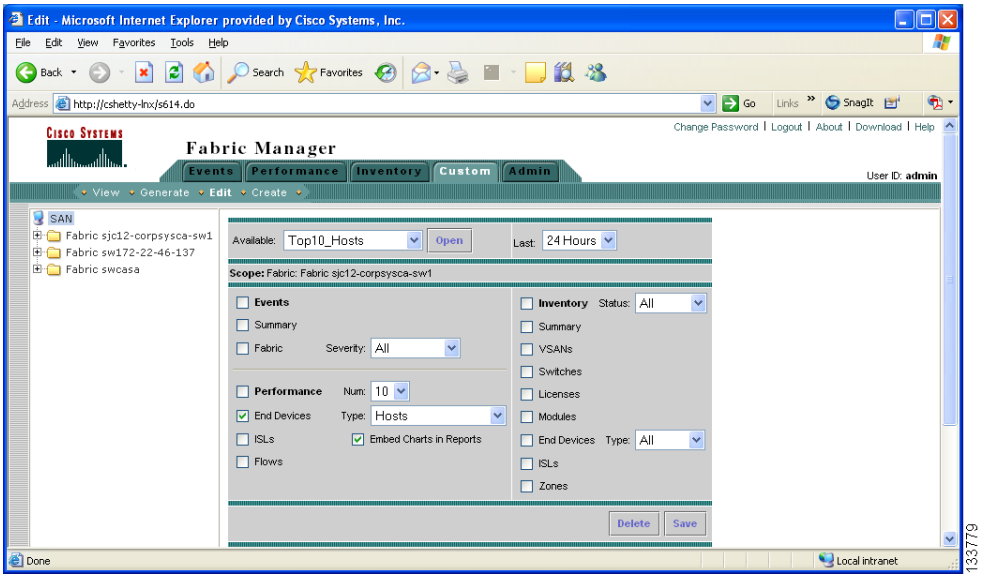

Step 3 Indicate the information you want in the report by checking the **Events**, **Performance**, or **Inventory** check boxes.

Step 4 Optionally, select a Severity for events, Status for inventory information, or Type of end device for performance information and inventory information.

Step 5 Click **Save** to save this report template.

Note You cannot change the SAN, fabric or VSAN the report is based on. Generate a new report for a new SAN, fabric or VSAN.

Admin

The Admin tab allows you to perform minor administrative and configuration tasks on the Fabric Manager Server sending data to your web client.

The Admin tab contains the following subtabs:

- Status—Displays the status of, and allows you to start and stop, the Database Server, Fabric Manager Server, and Performance Collector services on your server. You should only need to restart services if something is not working properly, or if too large a percentage of system resources are being consumed.
- Configure—Allows you to configure various parameters for Fabric Manager Server.
- Logs—Allows you to view all the logs from the various services running on the Fabric Manager Server.

Note If you see a database file lock error in the database log, you can fix it by shutting down and restarting the database server using the web client.

Starting, Restarting, and Stopping Services

To start, restart, or stop services using Fabric Manager Web Services, follow these steps:

Step 1 Choose the **Admin** tab followed by the **Status** tab.

You see a table of services and the status of each shown in Figure 6-25.

Figure 6-25 Fabric Manager Services and Their Status

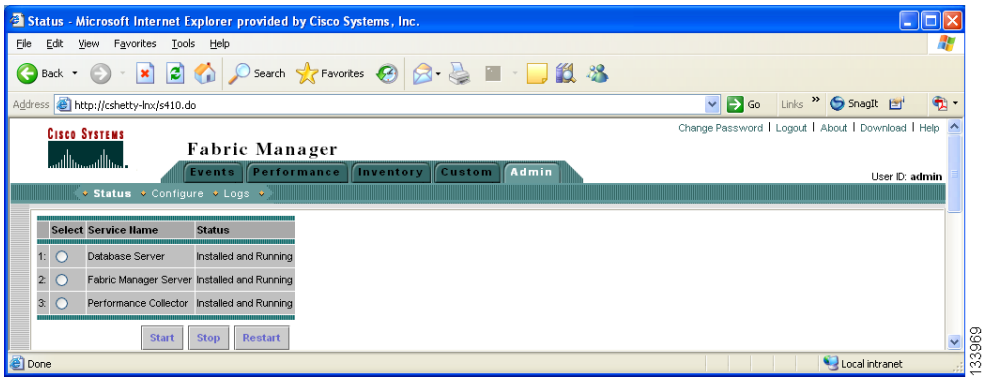

Step 2 Select the service(s) you want to start, restart, or stop.

Step 3 Click **Start**, **Stop,** or **Restart.**

The selected services are started, restarted, or stopped.

Note If the word “more” is in the Status column, you can click it to view a detailed status of the service.

Adding, Editing, and Removing Monitored Fabrics

Fabric Manager Web Services reports information gathered by Fabric Manager Server on any fabric known to the Fabric Manager Server.

To start monitoring a fabric from Fabric Manager Server using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab.

Step 2 Select **Fabrics** in the left navigation pane.

You see the list of fabrics (if any) monitored by Fabric Manager Server shown in Figure 6-26.

Figure 6-26 List of Fabrics Monitored by Fabric Manager Server

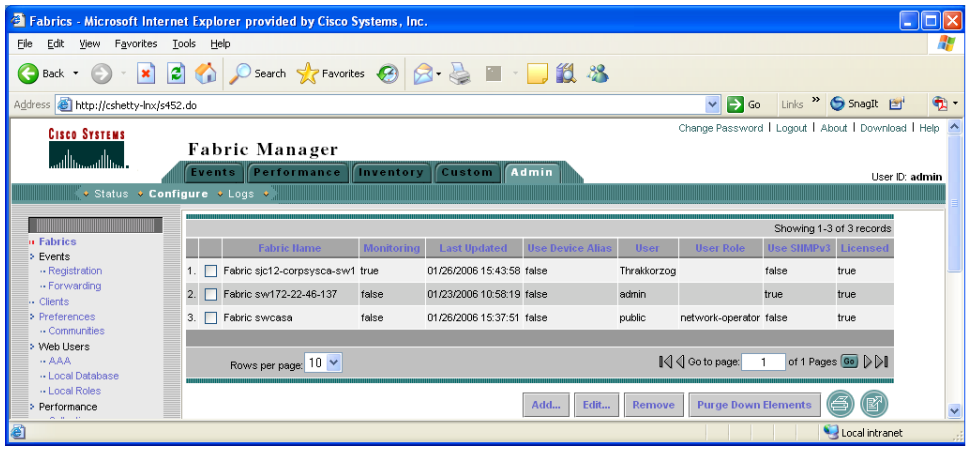

Step 3 Click **Add.**

You see the Add Fabric dialog box.

Step 4 Enter the seed switch IP address, read community, and write community for this fabric, and optionally check the SNMPV3 and Encryption check boxes. If you check SNMPV3, the fields Read Community, and Write Community change to Username and Password; provide them.

Step 5 Click **Add** to begin monitoring this fabric.

Step 6 Click **Close** to return to the Fabrics table.

To stop monitoring a fabric from Fabric Manager Server using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab.

You see the Configuration options.

Step 2 Click **Fabrics** in the left navigation pane.

You see the list of fabrics monitored by Fabric Manager Server shown in Figure 6-27.

Figure 6-27 List of Fabrics Monitored by Fabric Manager Server

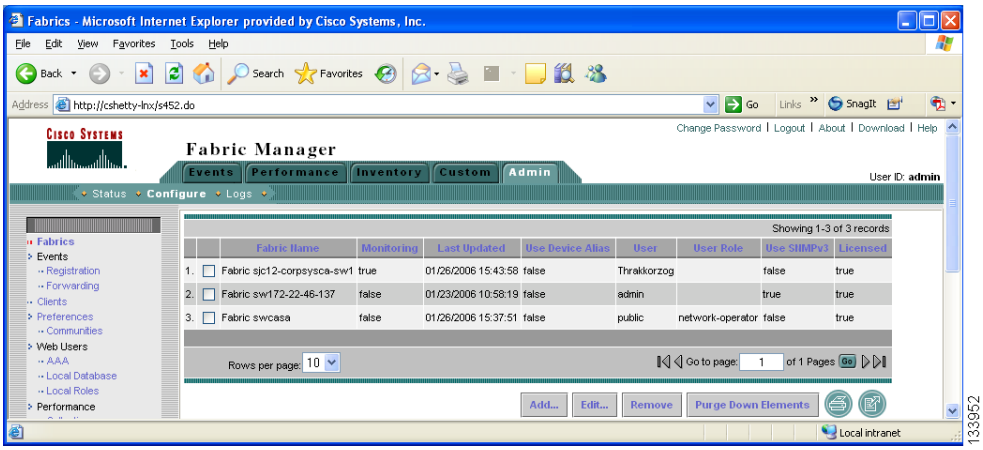

Step 3 Select a fabric and click **Remove** to discontinue data collection for that fabric (see Figure 6-27).

To edit a fabric from Fabric Manager Server using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab.

You see the Configuration options.

Step 2 Click **Fabrics** in the left navigation pane.

You see a list of fabrics monitored by Fabric Manager Server (see Figure 6-28).

Figure 6-28 List of Fabrics Monitored by Fabric Manager Server

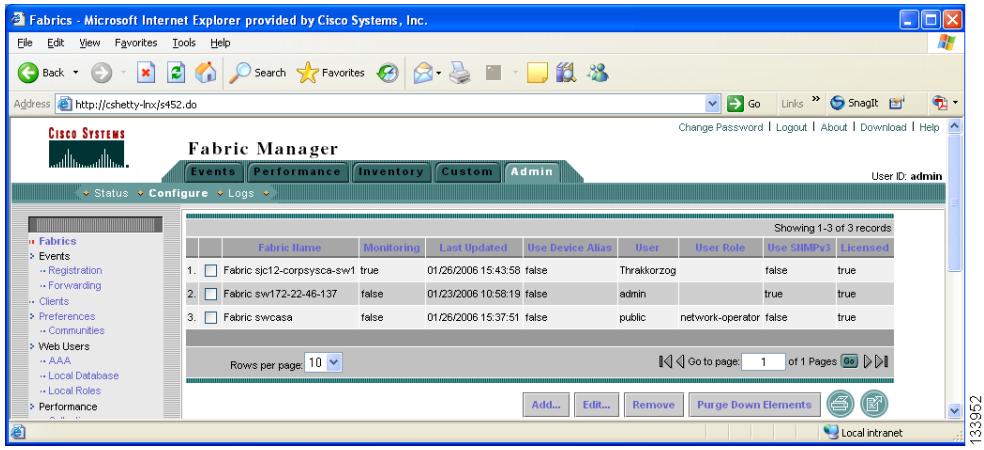

Step 3 Select a fabric and click **Edit** (see Figure 6-28).

You see the edit fabric dialog box.

Step 4 Enter a new fabric name, or change the monitoring status for this fabric.

Step 5 Click **Modify** to make the changes or click **Close**.

Viewing Trap and Syslog Registration Information

To view trap and syslog registration information from Fabric Manager Server using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab.

Step 2 Select **Registration** under Events.

You see the Registration screen shown in Figure 6-29.

Figure 6-29 Registration Screen

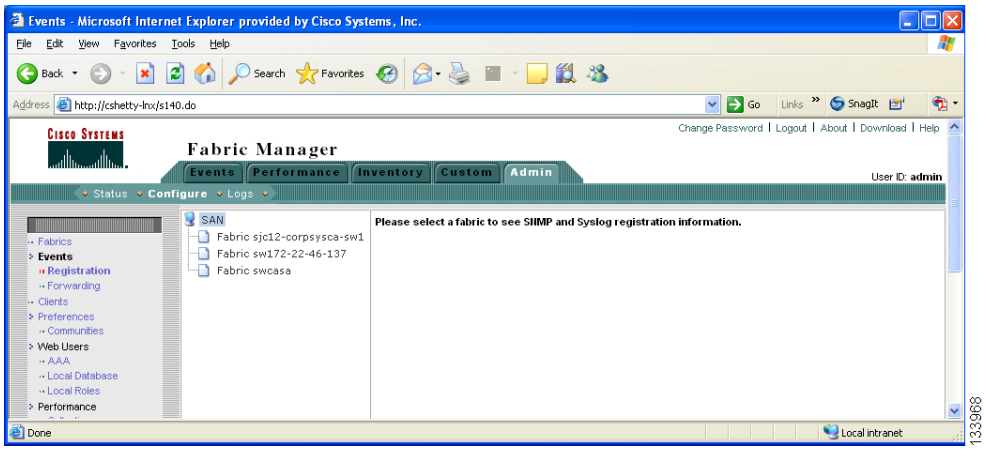

Step 3 Select a fabric to display registration information for that fabric.

Step 4 Optionally, click the **Print** icon or **Export Report** icon for a copy of the information.

Configuring Forwarding of Notifications for Events

You can use Fabric Manager Web Services to add and remove notification forwards for system messages.

To add a notification forward using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab, then select **Forwarding** under Events.

Step 2 Click **Add**.

You see the Add Notification dialog box shown in Figure 6-30.

Figure 6-30 Add Notification Dialog Box

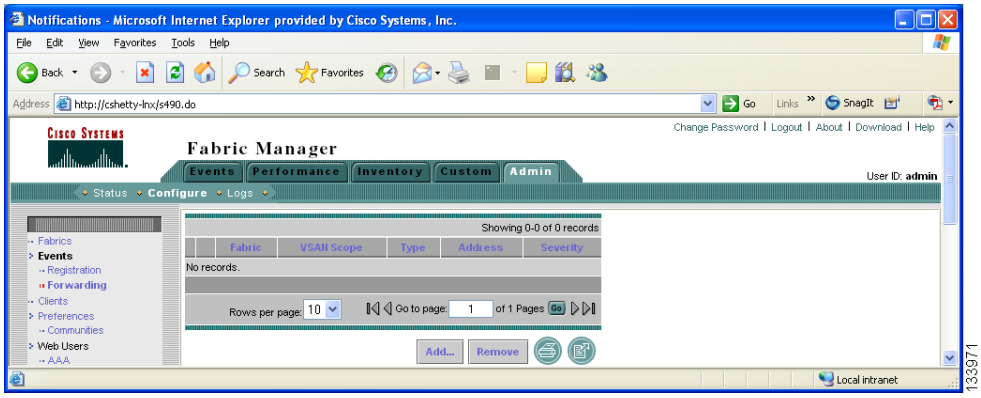

Step 3 Choose a Type of notification, either **E-Mail** or **Threshold Trap**. If you choose **Threshold Trap,** a Port field is added to the dialog.

Step 4 Choose the fabric for notification from the Fabric drop-down list.

Step 5 Either check the **VSAN Scope** check box to receive notifications for all VSANs, or enter the VSAN IDs in the ID List field to limit the VSANs for which you want to receive notifications.

Step 6 Enter the e-mail address for notifications in the Address field.

Step 7 Select the severity level of the messages to receive from the Severity drop-down list.

Step 8 Click **Add** to add the notification or click **Close**.

Note Click a column head to sort the forwarding information by that column.

To remove a notification forward using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab, then select **Forwarding** under Events.

You see the Forwarding screen.

Step 2 Check the check box in front of the notification forward to remove.

Step 3 Click **Remove**.

Note Click a column head to sort the forwarding information by that column.

Viewing and Disconnecting Clients

To view or disconnect clients from the Fabric Manager Server using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab, then select **Clients** in the left pane.

You see the Clients screen, showing a list of clients currently connected to the Fabric Manager Server (see Figure 6-31).

Figure 6-31 List of Clients Currently Connected to the Fabric Manager Server

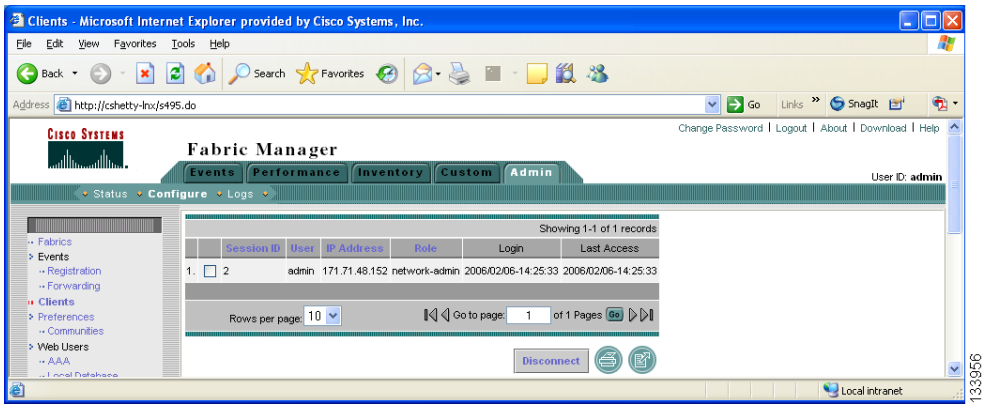

Step 2 Check the check box in front of a client to disconnect.

Step 3 Click **Disconnect.**

Note Click a column head to sort the client information by that column.

Configuring Fabric Manager Server Preferences

To configure Fabric Manager Server preferences, choose **Admin > Configure > Preferences** and follow the instructions on the screen.

Adding and Removing Communities

You can use Fabric Manager Web Services to add and remove communities.

To add a community fabric using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab, then select **Communities** under Preferences.

You see the Communities screen.

Figure 6-32 Communities Screen

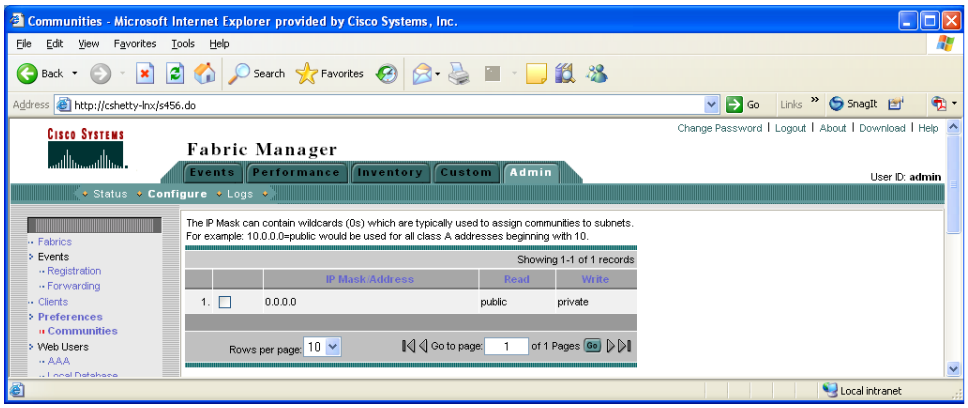

Step 2 Click **Add**.

You see the Add Fabric dialog box.

Step 3 Enter the IP mask or address of the community in the IP Mask/Address field.

Note The IP mask can contain wildcards (0s) you can use to assign communities to subnets.

Step 4 Enter the name of the community in the Community field.

Step 5 Click **Add** to add the fabric.

Note Click a column head to sort the community information by that column.

To remove a community using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab, then select **Communities** under Preferences.

You see the Communities screen.

Figure 6-33 Communities Screen

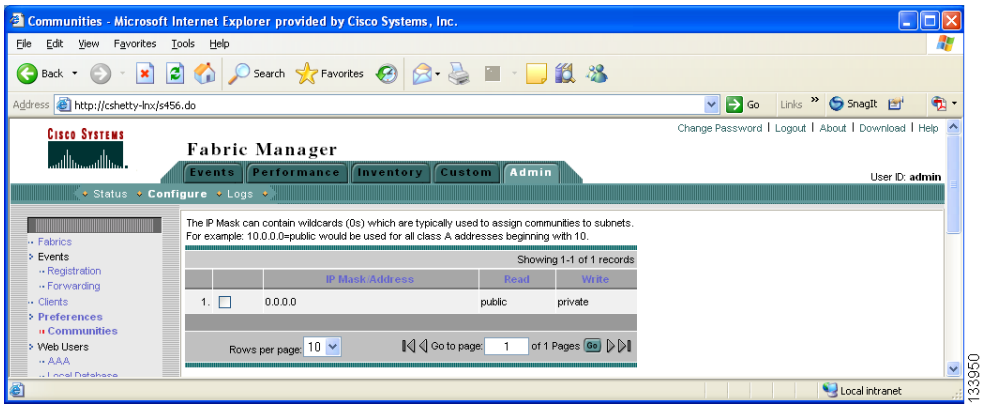

Step 2 Click the check box in front of the community to remove.

Step 3 Click **Remove**.

Note Click a column head to sort the community information by that column.

Note Cisco Fabric Manager 3.0(1) does not require you to make changes to the communities.properties file even if you are using a Cisco MDS 9020 switch or any third party devices.

Configuring AAA Information

To configure Fabric Manager Server preferences, click **Admin > Configure > Web Users > AAA** and follow the instructions on the screen.

Adding and Removing Users

You can use Fabric Manager Web Services to add and remove Web Services users.

To add a user using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab, then select **Local Database** under Web Users.

You see the Local Database screen shown in Figure 6-34.

Figure 6-34 Local Database Screen

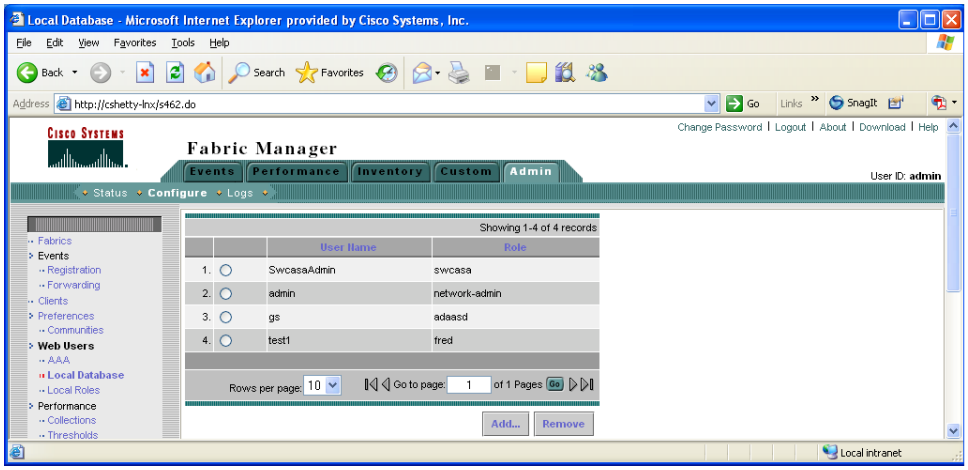

Step 2 Click **Add.**

You see the Add User dialog box.

Step 3 Enter the user name in the User Name field.

Note The user name **guest** is a reserved name (case insensitive). The guest user can only view reports. The guest user cannot change the guest password, nor can the guest user access the Admin tab in Fabric Manager Web Services.

Step 4 Select a role for the user from the Role drop-down list.

Step 5 Enter the password in the Password field.

Step 6 Enter the password again in the Confirm Password field.

Step 7 Click **Add** to add the user to the database.

Step 8 Repeat Steps 3 through 7 to continue adding users, or click **Close** to return to the Local Database screen.

Note Click a column head to sort the users by name or role.

To remove a user using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab, then select **Local Database** under Web Users. You see the Local Database screen shown in Figure 6-35.

Figure 6-35 Local Database Screen

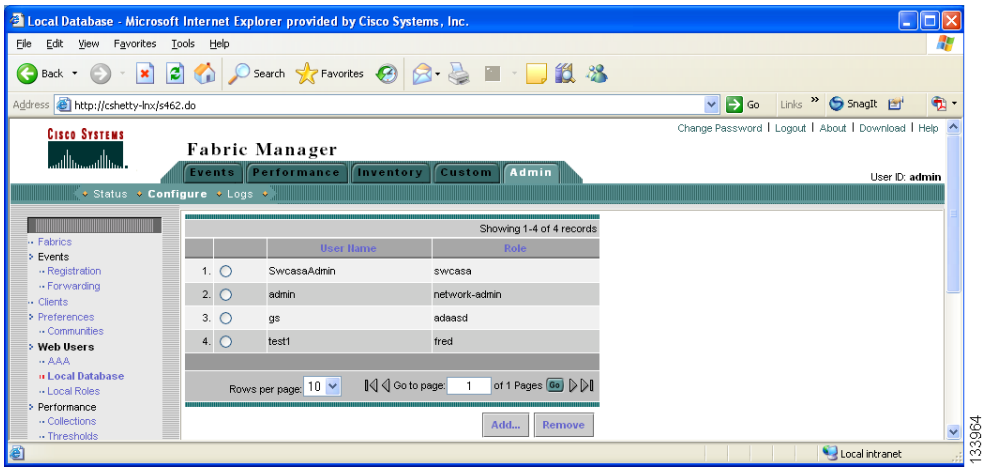

Step 2 Click the radio button in front of a user to remove.

Step 3 Click **Remove**.

Note Click a column head to sort the users by name or role.

Adding and Removing Roles

You can use Fabric Manager Web Services to add and remove Web Services roles.

To add a role using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab, then select **Local Roles** under Web Users.

You see the Local Roles screen shown in Figure 6-36.

Figure 6-36 Local Roles Screen

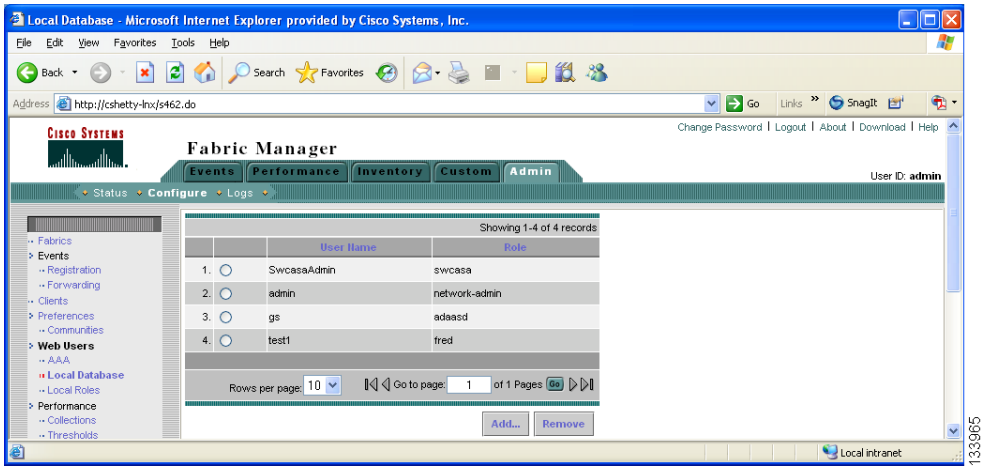

Step 2 Click **Add**.

You see the Add Roles dialog box.

Step 3 Enter the role name In the Role Name field.

Step 4 Select fabrics that the role can access from the Available Fabrics column and add them to the Selected Fabrics column.

Step 5 Click **Add** to add the role to the database.

Step 6 Repeat Steps 3 through 5 to continue adding roles, or click **Close** to return to the Local Roles screen.

Note Click a column head to sort the roles by role name or access.

To remove a role using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab, then select **Local Roles** under Web Users.

You see the Local Roles screen shown in Figure 6-37.

Figure 6-37 Local Roles Screen

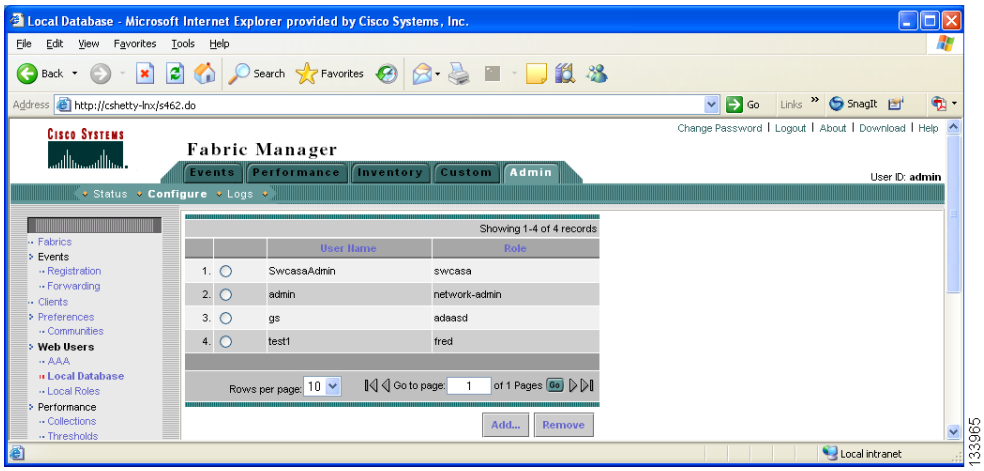

Step 2 Click the radio button in front of a role to remove.

Step 3 Click **Remove**.

Note Click a column head to sort the roles by role name or access.

Creating Performance Collections

If you are managing your fabrics with Performance Manager, you need to set up an initial set of flows and collections on the fabric. You can use Fabric Manager Web Services to add and remove performance collections.

Note You cannot manage performance collections for multiple devices through a single port interface. Since only one set of statistics exists per interface, Fabric Manager Web Services can manage performance collections for only one visible FL or iSCSI device through an interface.

To add a collection using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab, then select **Collections** under Performance.

You see the Collections screen.

Step 2 Click **Add**.

You see the Add Collection dialog box shown in Figure 6-38.

Figure 6-38 Add or Remove Collections Dialog Box

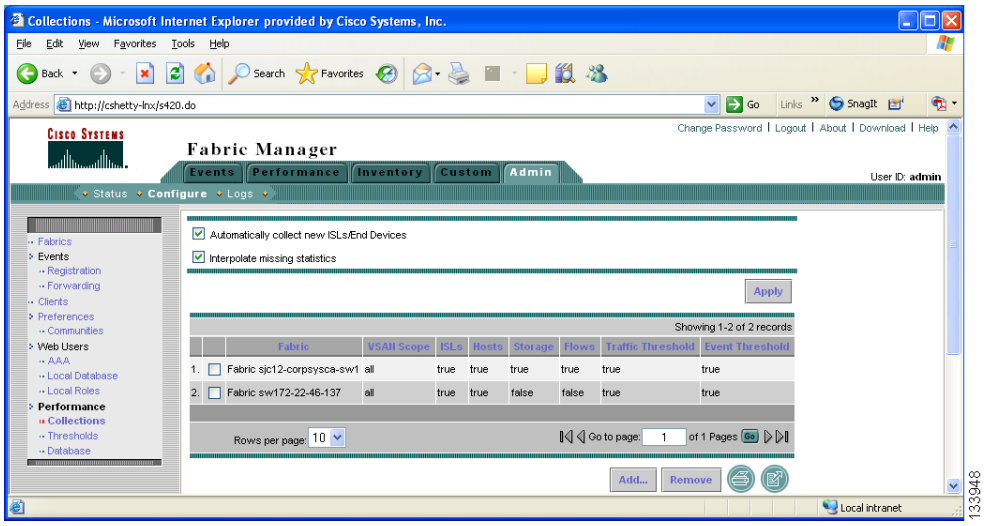

Step 3 Select a fabric for which to collect performance data from the Fabric drop-down list.

Step 4 Either check the **VSAN Scope** ‘check box to receive notifications for all VSANs, or enter the VSAN IDs in the ID List field to limit the VSANs for which you want to collect performance data.

Step 5 Check the check boxes for the type(s) of entities for which you want to collect performance data.

Step 6 Check the check boxes for the type(s) of thresholds you want to enable.

Step 7 Click **Create** to add the collection and add it to the table.

Step 8 Repeat Steps 3 through 8 to continue adding roles, or click **Close** to return to the Collections screen.

Note Click a column head to sort performance collections by that column.

To remove a collection using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab, then select **Collections** under Performance.

You see the Collections screen shown in Figure 6-39.

Figure 6-39 Add or Remove Collections Dialog Box

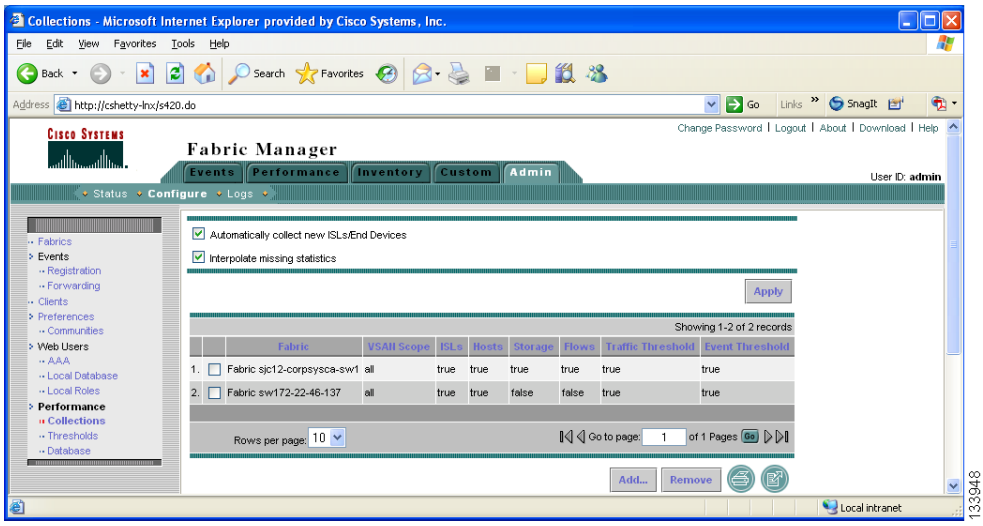

Step 2 Click the check box in front of a collection to remove.

Step 3 Click **Remove**.

Note Click a column head to sort performance collections by that column.

Configuring Collection Thresholds

To configure collection thresholds using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab, then select **Thresholds** under Performance.

You see the Thresholds screen shown in Figure 6-40.

Figure 6-40 Thresholds Screen

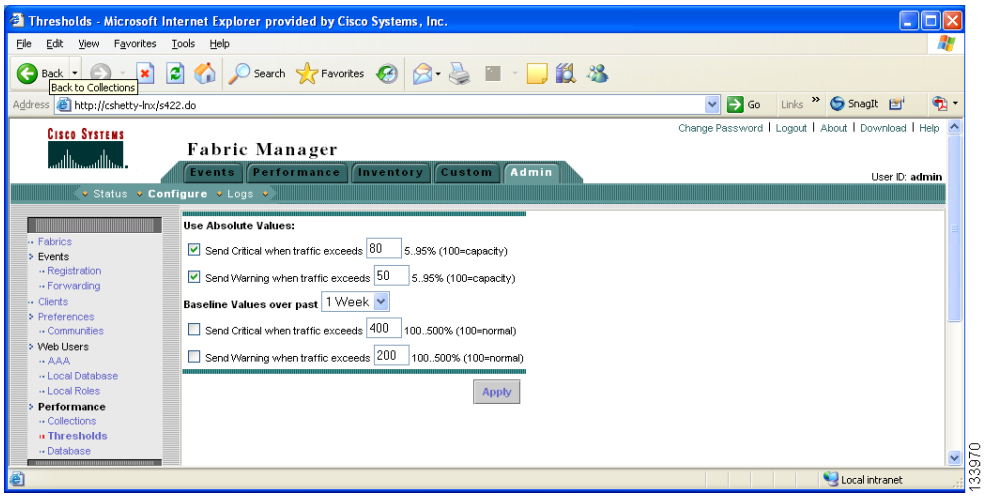

Step 2 If you are using absolute values, follow these steps. Otherwise skip to Step 3.

a. To configure conditions for sending Critical notifications, check the **Send Critical** check box. In the "...when traffic exceeds" field, enter a number (from 5 to 95) to indicate the percentage at which the Critical notification is sent. For example, entering **10** causes a notification to be sent when traffic at any given time exceeds 10% of capacity.

b. To configure conditions for sending Warning notifications, check the **Send Warning** check box. In the "...when traffic exceeds" field, enter a number (from 5 to 95) to indicate the percentage at which the Warning notification is sent. For example, entering **9** causes a notification to be sent when traffic at any given time exceeds 9% of capacity.

Step 3 Select the time period for the collection (**1 Week**, **1 Month**, or **1 Year**) from the Baseline Values over past drop-down list. The baseline value represents the sum of the absolute values.

a. To configure conditions for sending Critical notifications, check the **Send Critical** check box. In the "...when traffic exceeds" field, enter a number to indicate the percentage at which the Critical notification is sent. For example, entering **300** causes a notification to be sent when traffic for the selected period exceeds 300% of capacity.

b. To configure conditions for sending Warning notifications, check the **Send Warning** check box. In the "...when traffic exceeds" field, enter a number to indicate the percentage at which the Warning notification is sent. For example, entering **150** causes a notification to be sent when traffic for the selected period exceeds 150% of capacity.

Step 4 Click **Apply**.

Configuring the RRD Database

To configure the RRD database using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Configure** tab, then select **Database** under Performance.

You see the Database (collection interval) screen shown in Figure 6-41.

Figure 6-41 Configure Database

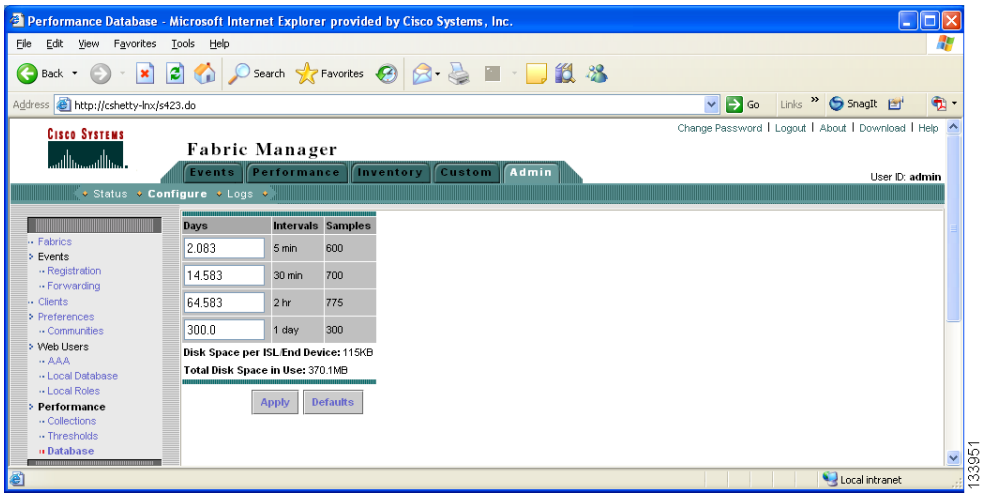

Step 2 Enter the number of days to collect samples at 5-minute intervals in the top row of the Days column.

Step 3 Enter the number of days to collect samples at 30-minute intervals in the second row of the Days column.

Step 4 Enter the number of days to collect samples at 2-hour intervals in the third row of the Days column.

Step 5 Enter the number of days to collect samples at 1-day intervals in the bottom row of the Days column.

Step 6 Click **Apply** to apply your changes, or click **Defaults** to reset the file sizes to the default values.

If you are applying new values, or if the current values are not the default values, you see a message indicating that conversion of the RRD files will take a certain amount of time and that the database will be unavailable until then. The time it takes depends on the difference between the old and new values.

Note The system allows only one conversion process at a time. Once you start the conversion, the Apply and Default buttons change to Refresh and Cancel so that another process cannot be inadvertently started. This display is the same for all browsers accessing this server during this time. Click **Refresh** to view the latest progress of the conversion. Click **Cancel** to cancel the conversion job. If the job is successfully canceled, you see the Apply and Default buttons again. If the cancel is not successful, you see a message indicating that the cancellation has failed.

If you want to perform this procedure, it is best to perform it before collecting a lot of data. Otherwise, the conversion can take a long time to complete.

Viewing Log Information

You may occasionally want to view logs such as the Fabric Manager Server log and the Performance Manager log. These processes have no corresponding GUI to allow you to view information about these log files. If you see errors, preserve these two files for viewing.

To view log information using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Admin** tab followed by the **Logs** tab.

You see a list of viewable logs in the left column.

Step 2 Click the log to view.

CHAPTER 7

Performance Manager

The primary purpose of Fabric Manager is to manage the network. A key management capability is network performance monitoring.

Performance Manager Architecture

Performance Manager gathers network device statistics historically and provides this information graphically using a web browser. It presents recent statistics in detail and older statistics in summary. Performance Manager also integrates with external tools such as Cisco Traffic Analyzer.

The Performance Manager has three operational stages:

- Definition—Uses two configuration wizards. The Flow Wizard sets up flows in the switches while the Collection Wizard creates a collection configuration file.
- Collection—Reads the configuration file and collects the desired information.
- Presentation—Generates web pages to present the collected data.

Performance Manager can collect statistics for ISLs, hosts, storage elements, and configured flows. Flows are defined based on a host-to-storage (or storage-to-host) link. Performance Manager gathers statistics from across the fabric based on collection configuration files. These files determine which SAN elements and SAN links Performance Manager gathers statistics for. Based on this configuration, Performance Manager communicates with the appropriate devices (switches, hosts, or storage elements) and collects the appropriate information at fixed five-minute intervals.

Performance Manager uses a round-robin database to hold the statistical data collected from the fabric. This data is stored based on the configured parameters in the collection configuration file. At each polling interval, Performance Manager gathers the relevant statistics and stores them in the round-robin database. This database is a fixed size and will not grow beyond its preset limits.

Performance Manager creates a series of archived data to hold summarized information present in the real-time round-robin database. This archived data is used to generate daily, weekly, monthly, and yearly consolidated reports. In this way, Performance Manager maintains significant historical data without the cost of an ever-increasing database size.

Data Interpolation

One of the unique features of Performance Manager is its ability to interpolate data when statistical polling results are missing or delayed. Other performance tools may store the missing data point as zero, but this can distort historical trending. Performance Manager interpolates the missing data point by comparing the data point that preceded the missing data and the data point stored in the polling interval after the missing data. This maintains the continuity of the performance information.

Data Collection

One year's worth of data for two variables (Rx and Tx bytes) requires a round-robin database (rrd) file size of 76K. If errors and discards are also collected, the rrd file size becomes 110K. The default internal values are:

- 600 samples of 5 minutes (2 days and 2 hours)
- 700 samples of 30 minutes (2 days and 2 hours, plus 12.5 days)
- 775 samples of 2 hours (above plus 50 days)
- 300 samples of 1 day (above plus300 days, rounded up to 365)

A 1000-port SAN requires 110 MB for a year's worth of historical data that includes errors and discards. If there were 20 switches in this SAN with equal distribution of fabric ports, about two to three SNMP packets per switch would be sent every 5 minutes for a total of about 100 request or response SNMP packets required to monitor the data.

Flows, because of their variable counter requests, are more difficult to predict storage space requirements for. But as a rule of thumb, each extra flow adds another 76 kB.

The Performance Manager collector runs as a background process on the various supported operating systems. On Microsoft Windows, it runs as a service.

Using Performance Thresholds

The Performance Manager Configuration Wizard allows you to set up two thresholds that will trigger events when the monitored traffic exceeds the percent utilization configured. These event triggers can be set as either Critical or Warning events that are reported on the Fabric Manager web client Events browser page.

Absolute value thresholds apply directly to the statistics gathered. These statistics, as a percent of the total link capacity, are compared to the percent utilization configured for the threshold type. If the statistics exceed either configured threshold, an event is shown on the Fabric Manager web client Events tab.

Baseline thresholds create a threshold that adapts to the typical traffic pattern for each link for the same time window each day, week, or every two weeks. Baseline thresholds are set as a percent of the average (110% to 500%), where 100% equals the calculated weighted average. Figure 7-1 shows an example of setting a baseline threshold for a weekly or daily option.

Figure 7-1 Baseline Threshold Example

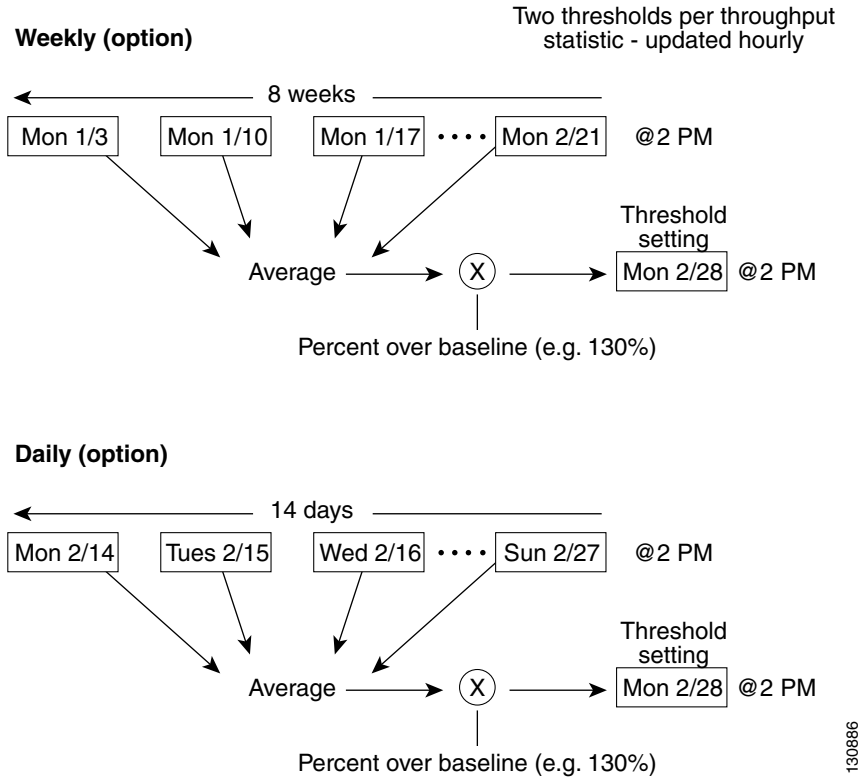

The threshold is set for Monday at 2 p.m. The baseline threshold is set at 130% of the average for that statistic. The average is calculated from the statistics value that occurred at 2 p.m. on Monday, for every prior Monday (for the weekly option) or the statistics value that occurred at 2 p.m. on each day, for every prior day (for the daily option).

Quick Data Collector and Flow Setup Wizards

The Performance Manager Flow and Performance Manager Setup wizards greatly simplify configuration. All you need to do is select the categories of statistics to capture and the wizards provide a list of flows and links to monitor. You can remove entries if desired, or just accept the provided list and start data collection. Statistics for host and storage links are not associated with a specific port on a switch, so you do not lose long term statistics if a connection is moved to a different port.

CHAPTER 8

Authentication in Fabric Manager

Fabric Manager contains interdependent software components that communicate with the switches in your fabric. These components use varying methods to authenticate to other components and switches. This chapter describes these authentication steps and the best practices for setting up your fabric and components for authentication.

This chapter contains the following sections:

Fabric Manager Authentication Overview

Fabric Manager contains multiple components that interact to manage a fabric. These components include:

- Fabric Manager Client
- Fabric Manager Server
- Performance Manager
- Interconnected fabric of Cisco MDS 9000 switches and storage devices
- AAA server (optional)

Figure 8-1 shows an example configuration for these components.

Figure 8-1 Fabric Manager Authentication Example

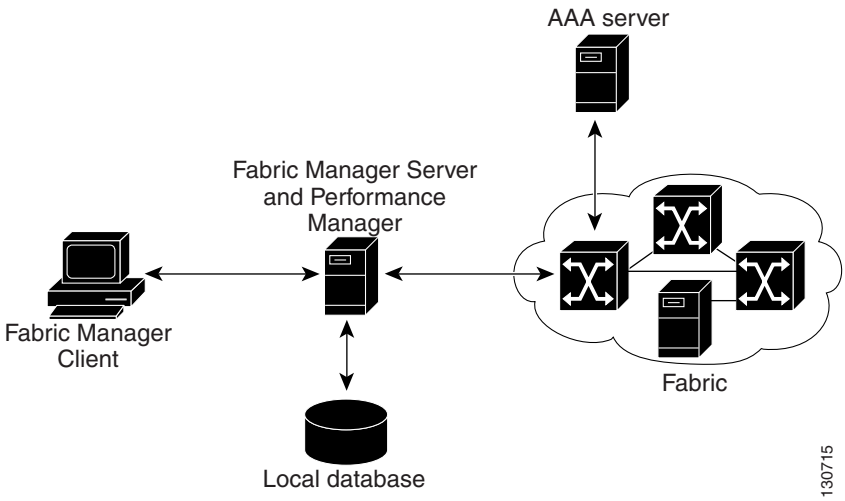

Administrators launch Fabric Manager Client and select the seed switch that is used to discover the fabric. The user name and password used are passed to Fabric Manager Server and used to authenticate to the seed switch. If this user name and password are not a recognized SNMP user name and password, either Fabric Manager Client or Fabric Manager Server opens a CLI session to the switch (SSH or Telnet) and retries the user name/password pair. If the user name and password are recognized by the switch in either the local switch authentication database or through a remote AAA server, then the switch creates a temporary SNMP user name that is used by Fabric Manager Client and server.

Note You may encounter a delay in authentication if you use a remote AAA server to authenticate Fabric Manager or Device Manager.

Note You must allow CLI sessions to pass through any firewall that exists between Fabric Manager Client and Fabric Manager Server. See the "Running Fabric Manager Behind a Firewall" section on page 2-25.

Note We recommend that you use the same password for the SNMPv3 user name authentication and privacy passwords as well as the matching CLI user name and password.

Best Practices for Discovering a Fabric

Fabric Manager Server monitors multiple physical fabrics under the same user interface. This facilitates managing redundant fabrics. A licensed Fabric Manager Server maintains up-to-date discovery information on all configured fabrics so device status and interconnections are immediately available when you launch Fabric Manager Client.

We recommend you use these best practices for discovering your network and setting up Performance Manager. This ensures that Fabric Manager Server has a complete view of the fabric. Subsequent Fabric Manager Client sessions can filter this complete view based on the privileges of the client logging in. For example, if you have multiple VSANs in your fabric and you create users that are limited to a subset of these VSANs, you want to initiate a fabric discovery through Fabric Manager Server using a network administrator or network operator role so that Fabric Manager Server has a view of all the VSANs in the fabric. When a VSAN-limited user launches Fabric Manager Client, that user sees only the VSANs they are allowed to manage.

We recommend you use these best practices for discovering your network and setting up Performance Manager.

Setting Up Discovery for a Fabric

To ensure that Fabric Manager Server discovers your complete fabric, follow these steps:

Step 1 Create a special Fabric Manager administrative user name in each switch on your fabric with network administrator or network operator roles. Or, create a special Fabric Manager administrative user name in your AAA server and set every switch in your fabric to use this AAA server for authentication.

Step 2 Verify that the roles used by this Fabric Manager administrative user name are the same on all switches in the fabric and that this role has access to all VSANs.

Step 3 Launch Fabric Manager Client using the Fabric Manager administrative user. This ensures that your fabric discovery includes all VSANs.

Step 4 Set Fabric Manager Server to continuously monitor the fabric. See the "Fabric Manager Server Fabric Monitoring and Removal" section on page 3-6.

Step 5 Repeat Step 4 for each fabric that you want to manage through Fabric Manager Server.

Performance Manager Authentication

Performance Manager uses the user name and password information stored in the Fabric Manager Server database. If this information changes on the switches in your fabric while Performance Manager is running, you need to update the Fabric Manager Server database and restart Performance Manager. Updating the Fabric Manager Server database requires removing the fabric from Fabric Manager Server and rediscovering the fabric.

To update the user name and password information used by Performance Manager, follow these steps:

Step 1 Click **Server > Admin** in Fabric Manager.

You see the Server Admin dialog box.

Step 2 Click the **Fabrics** tab to view the fabrics currently monitored by Fabric Manager Server.

You see the Server Admin dialog box shown in Figure 8-2.

Figure 8-2 Server Admin Dialog Box

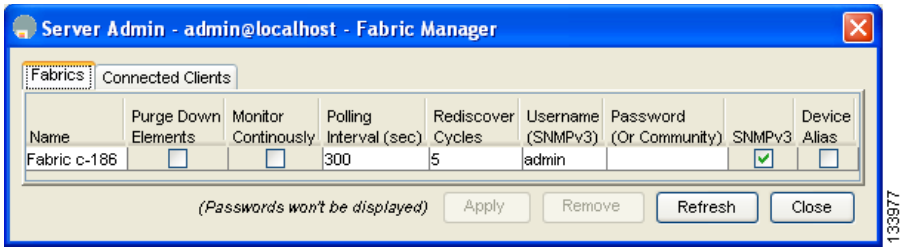

Step 3 Click the fabrics that have updated user name and password information.

Step 4 Click **Remove** to remove these fabrics from Fabric Manager Server.

Step 5 Choose **File > Open Fabric.**

You see the Open Fabric dialog box shown in Figure 8-3.

Figure 8-3 Open Fabric Dialog Box

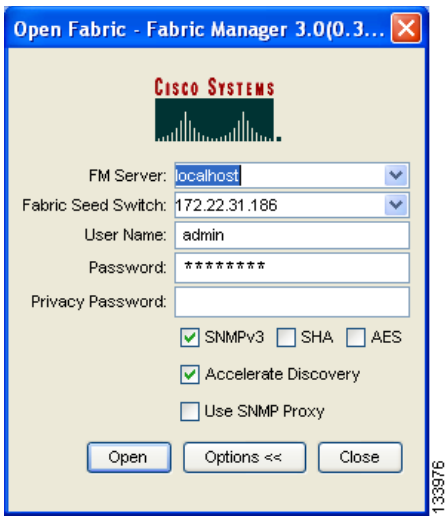

Step 6 Set the seed switch and the appropriate user name and password to rediscover the fabric.

Step 7 Click **Open** to rediscover the fabric. Fabric Manager Server updates its user name and password information (see Figure 8-3).

Step 8 Repeat Step 5 through Step 7 for any fabric that you need to rediscover.

Step 9 Click **Performance > Collector > Restart** to restart Performance Manager and use the new user name and password.

Fabric Manager Web Services Authentication

Fabric Manager Web Services does not communicate directly with any switches in the fabric. Fabric Manager Web Services uses its own user name and password combination that is either stored locally or stored remotely on an AAA server.

We recommend that you use a RADIUS or TACACS+ server to authenticate users in Fabric Manager Web Services.

To configure Fabric Manager Web Services to use RADIUS authentication, follow these steps:

Step 1 Launch Fabric Manager Web Services. See the "Launching Fabric Manager Web Services" section on page 6-5.

Step 2 Click the **Admin** tab **> Web Users** to update the authentication used by Fabric Manager Web Services.

Step 3 Click **AAA**.

Step 4 Set the authenticationmode attribute to **radius**.

Step 5 Set the RADIUS server name, shared secret, authentication method, and ports used for up to three RADIUS servers.

Step 6 Click **Modify** to save this information.

To configure Fabric Manager Web Services to use TACACS+ authentication, follow these steps:

Step 1 Launch Fabric Manager Web Services. See the "Launching Fabric Manager Web Services" section on page 6-5.

Step 2 Click **Admin > Web Users** to update the authentication used by Fabric Manager Web Services.

Step 3 Click **AAA**.

Step 4 Set the authenticationmode attribute to **tacacs**.

Step 5 Set the TACACS+ server name, shared secret, authentication method, and port used for up to three TACACS+ servers.

Step 6 Click **Modify** to save this information.

CHAPTER 9

Cisco Traffic Analyzer

Cisco Traffic Analyzer is a version of network top (ntop) software that is modified to support Fibre Channel and SCSI.

This chapter contains the following sections:

Using Cisco Traffic Analyzer with Performance Manager

Performance Manager works in conjunction with Cisco Traffic Analyzer to monitor and manage the traffic on your fabric. Using Cisco Traffic Analyzer with Performance Manager requires the following components:

- A configured Fibre Channel Switched Port Analyzer (SPAN) destination (SD) port to forward Fibre Channel traffic.
- A Port Analyzer Adapter 2 (PAA-2) to convert the Fibre Channel traffic to Ethernet traffic.
- Cisco Traffic Analyzer software to analyze the traffic from the PAA-2.

Figure 9-1 shows how Performance Manager works with Cisco Traffic Analyzer to monitor traffic on your fabric.

Figure 9-1 Overview of Performance Manager Working with Cisco Traffic Analyzer

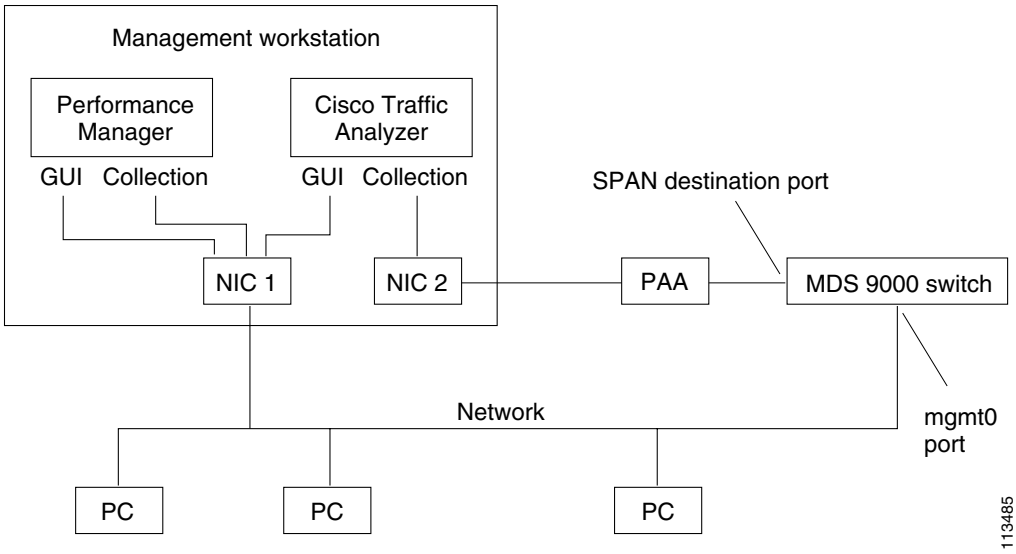

Understanding SPAN

The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic though a Fibre Channel interface. Traffic through any Fibre Channel interface can be replicated to a special port called the SPAN destination port (SD port). Any Fibre Channel port in a switch can be configured as an SD port. Once an interface is in SD port mode, it cannot be used for normal data traffic. You can attach a Fibre Channel analyzer to the SD port to monitor SPAN traffic.

SD ports do not receive frames, they merely transmit a copy of the SPAN source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports (see Figure 9-2).

Figure 9-2 SPAN Transmission

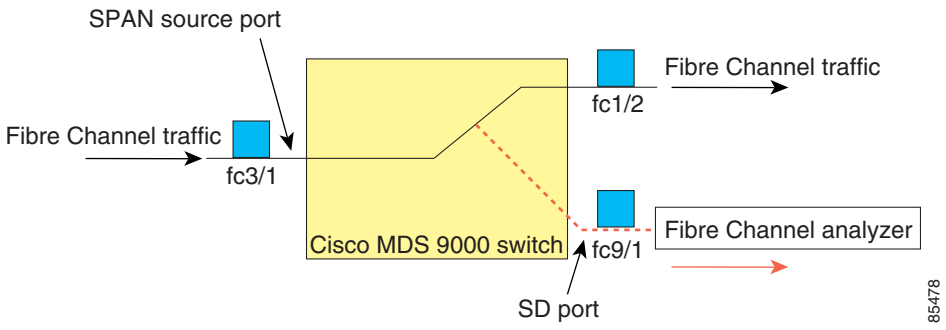

For information on configuring SPAN, refer to the *Cisco MDS 9000 Family CLI Configuration Guide.*

Understanding the PAA-2

The PAA-2 enables effective, low-cost analysis of Fibre Channel traffic. The device is a standalone Fibre Channel-to-Ethernet adapter, designed primarily to analyze SPAN traffic from a Fibre Channel port on a Cisco MDS 9000 Family switch. The main function of the Port Analyzer Adapter 2 is to encapsulate Fibre Channel frames into Ethernet frames. This allows low-cost analysis of Fibre Channel traffic while leveraging the existing Ethernet infrastructure.

The PAA-2 allows you to examine Fibre Channel frames of various sizes. Fibre Channel frames from Layers 2, 3, and 4 may be examined without network disruption.

Understanding Cisco Traffic Analyzer

Performance Manager collects Fibre Channel level performance statistics using SNMP to access counters on Cisco MDS 9000 Family switches. To view detailed SCSI I/O statistics, you need to look at the data on an SD port with the help of Cisco Traffic Analyzer, which uses the Cisco Port Analyzer Adapter 2 (PAA-2).

Cisco Traffic Analyzer provides real-time analysis of SPAN traffic or analysis of captured traffic through a Web browser user interface. Traffic encapsulated by one or more Port Analyzer Adapter 2 products can be analyzed concurrently with a single workstation running Cisco Traffic Analyzer, which is based on ntop, a public domain software enhanced by Cisco for Fibre Channel traffic analysis.

Round-trip response times, SCSI I/Os per second, SCSI read or traffic throughput and frame counts, SCSI session status, and management task information are monitored. Additional statistics are also available on Fibre Channel frame sizes and network management protocols.

For seamless performance analysis and troubleshooting, Cisco Traffic Analyzer can be launched in-context from Fabric Manager. Port world wide name (pWWN), Fibre Channel ID (FC ID), FC alias, and VSAN names are passed to Cisco Traffic Analyzer.

Cisco Traffic Analyzer must be downloaded and installed separately from the following website:

http://www.cisco.com/kobayashi/sw-center/sw-stornet.shtml.

Cisco Traffic Analyzer software is available under the Port Analyzer Adapter link. See the "Installing Cisco Traffic Analyzer" section on page 9-4.

Caution Cisco Traffic Analyzer for Fibre Channel throughput values are not accurate when used with the original Cisco Port Analyzer Adapter (PAA) if data truncation is enabled. PAA Version 2 (product ID DS-PAA_2) is required to achieve accurate results with truncation, because it adds a count that enables Cisco Traffic Analyzer to determine how many data bytes were actually transferred.

Note Refer to the *Cisco MDS 9000 Family CLI Configuration Guide* for information on configuring the settings for your SPAN destination port. It is important that the data you collect through this port matches the data collected by Performance Manager through the mgmt0 port. If the data does not match, you cannot view Cisco Traffic Analyzer information through a Traffic Analyzer link on the detail page of a Performance Manager report.

Installing Cisco Traffic Analyzer

To install Cisco Traffic Analyzer on a UNIX workstation, follow these steps:

Step 1 Open a browser and go to the following website to access the web page where Cisco Traffic Analyzer is available:
http://cisco.com/cgi-bin/tablebuild.pl/mds-fm.

Step 2 Download fc-ntop.tar.gz and install it using the instructions at the following website:
http://www.ntop.org.

Step 3 Verify that the Fibre Channel port on the PAA-2 is connected to the SD port on the switch (Figure 9-1).

Step 4 Verify that the Ethernet port on the PAA-2 is connected to the workstation running Cisco Traffic Analyzer.

Step 5 Click **Interfaces > SPAN** in Device Manager to configure SPAN on the required switch ports.

Step 6 Click **Interfaces > SPAN** in Device Manager to verify that the Fibre Channel port connected to the PAA-2 is configured as an SD port. The port mode of the destination interface must be SD.

Step 7 Click the **Sessions** tab in Device Manager to verify the correct destination and source of traffic (ingress).

Caution Cisco Traffic Analyzer must not be used with the PAA-2 in Management mode (MNM). Refer to the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*.

To install Cisco Traffic Analyzer on a Windows workstation, follow these steps:

Step 1 Open a browser and go to the following website to access the web page where Cisco Traffic Analyzer is available:
http://cisco.com/cgi-bin/tablebuild.pl/mds-fm.

Step 2 Download ntop-win32.zip and save it on your workstation.

Step 3 Unzip the downloaded file.

Note You need the WinPcap library file to use Cisco Traffic Analyzer on a Microsoft Windows system.You can download this file from the following website:
http://winpcap.polito.it/.

Step 4 Open a command prompt and change directories to your ntop installation directory.

Step 5 Type **ntop -i** or install ntop as a service on Windows by following these steps:

a. Type **ntop /i** to install ntop as a service.

b. Choose **Start > Programs > Administrative Tools > Services** to access the Windows Services Panel.

c. Right-click **ntop** and choose **properties**. You see the Properties dialog box.

d. Set the Start Parameters to **-i *interface number***, where *interface number* is the number of the interface on your workstation that connects to the PAA-2.

e. Click **Start** to start ntop on that interface.

Note Subsequent restarts of the ntop service do not require setting the -i option. unless you are changing the interface that connects to the PAA-2.

Step 6 Verify that the Fibre Channel port on the PAA-2 is connected to the SD port on the switch (Figure 9-1).

Step 7 Verify that the Ethernet port on the PAA-2 is connected to the workstation running Cisco Traffic Analyzer.

Step 8 Click **Interfaces > SPAN** in Device Manager to configure SPAN on the required switch ports.

Step 9 Click the **Sources** tab in Device Manager to verify that the Fibre Channel port connected to the PAA-2 is configured as an SD port. The port mode of the destination interface must be SD.

Step 10 Click the **Sessions** tab in Device Manager to verify the correct destination and source of traffic (ingress).

Tip To modify the script that launches ntop (ntop.sh or ntop.bat), follow the instructions provided within the script file. Create a backup of the original script before modifying the file.
—Linux platforms use the shell script path. The ntop output is sent to the syslog file (/var/log/messages by default).
—Windows platforms use the batch file. The ntop output is sent to a file located in the same directory as the one from which ntop is launched.

Configuring Performance Manager for Use with Cisco Traffic Analyzer

Fabric Manager Release 2.1(2) or later supports Cisco Traffic Analyzer directly from Fabric Manager Web Services.

To configure Performance Manager to work with Cisco Traffic Analyzer for Fabric Manager releases prior to Release 2.1(2), follow these steps:

Step 1 Get the following three pieces of information:

- The IP address of the management workstation on which you are running Performance Manager and Cisco Traffic Analyzer.
- The path to the directory where Cisco Traffic Analyzer is installed.
- The port that is used by Cisco Traffic Analyzer (the default is 3000).

Step 2 Start Cisco Traffic Analyzer.

a. Choose **Performance > Traffic Analyzer > Open**.

b. Enter the URL for Cisco Traffic Analyzer, in the format

```
http://ip address:port number
```

where:

ip address is the address of the management workstation on which you have installed the Cisco Traffic Analyzer, and :*port number* is the port that is used by Cisco Traffic Analyzer (the default is :3000).

c. Click **OK**.

d. Choose **Performance > Traffic Analyzer > Start**.

e. Enter the location of Cisco Traffic Analyzer, in the format

```
D:\directory\ntop.bat
```

where:

D: is the drive letter for the disk drive where Cisco Traffic Analyzer is installed, and

directory is the directory containing the ntop.bat file.

f. Click **OK**.

Step 3 Create the flows you want Performance Manager to monitor, using the Flow Configuration Wizard.

Step 4 Define the data collection you want Performance Manager to gather, using the Performance Manager Configuration Wizard.

a. Choose the VSAN you want to collect information for or choose All VSANs.

b. Check the types of items you want to collect information for (hosts, ISLs, storage devices, and flows).

c. Enter the URL for Cisco Traffic Analyzer in the format

```
http://ip address/directory
```

where:

ip address is the address of the management workstation on which you have installed Cisco Traffic Analyzer, and *directory* is the path to the directory where Cisco Traffic Analyzer is installed.

d. Click **Next**.

e. Review the end devices and links you selected to make sure this is the data you want to collect.

f. Click **Finish** to begin collecting data.

Note Data is not collected for JBOD or for virtual ports. If you change the data collection configuration parameters during a data collection, you must stop and restart the collection process for your changes to take effect.

Step 5 Click **Cisco Traffic Analyzer** at the top of the Host or Storage detail pages to view Cisco Traffic Analyzer information, or choose **Performance > Traffic Analyzer > Open**. Cisco Traffic Analyzer will not open unless ntop has been started already.

Note For information on capturing a SPAN session and starting a Cisco Traffic Analyzer session to view it, refer to the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note.*

Note For information on viewing and interpreting your Performance Manager data, see the "Historical Performance Monitoring" section on page 54-4.

For information on viewing and interpreting your Cisco Traffic Analyzer data, refer to the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*.

For performance drill-down, Fabric Manager Server can launch Cisco Traffic Analyzer in context from the Performance Manager graphs. The aliases associated with hosts, storage devices, and VSANs are passed to Cisco Traffic Analyzer to provide consistent, easy identification.

Accessing Traffic Analyzer from Fabric Manager Web Services

Fabric Manager Release 2.1(2) or later supports discovering instances of Traffic Analyzer and SPAN ports configured within your fabric.

Fabric Manager Web Services supports the following Traffic Analyzer integration features:

- SCSI I/O Traffic Analyzer pages can be viewed within the Web client.
- Traffic Analyzer can reside on a different server than Performance Manager.
- Performance Manager integrates with multiple servers running Traffic Analyzer.
- Instances of Traffic Analyzer servers can be discovered by Fabric Manager Server.
- Web client report lists SPAN destination ports and associations with Traffic Analyzers.

To access an instance of Traffic Analyzer running in your fabric using Fabric Manager Web Services, follow these steps:

Step 1 Choose the **Performance tab** then the **Traffic Analyzer tab**.

You see a summary table of all SPAN destination ports and configured Traffic Analyzers in your fabric (see Figure 9-3). The source column shows the ports that are monitored by the SPAN destination port.

Figure 9-3 Traffic Analyzer in Fabric Manager Web Services

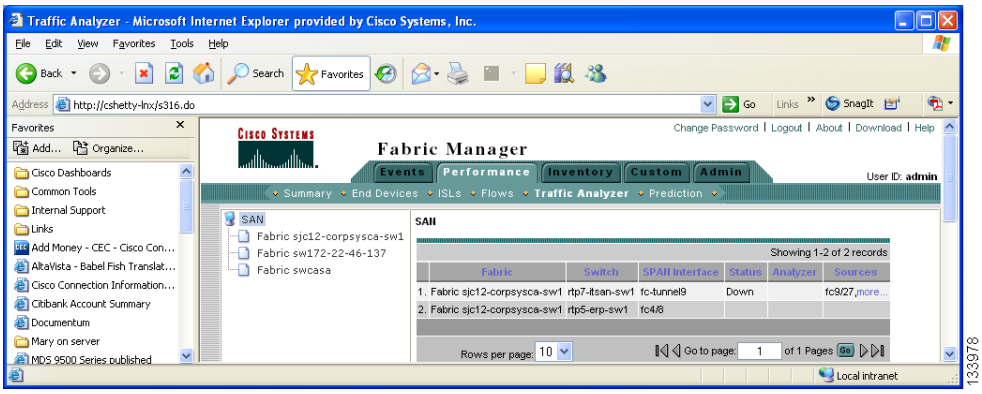

Step 2 Click a Traffic Analyzer to launch that Traffic Analyzer within Fabric Manager Web Services.

To rediscover instances of Traffic Analyzer running in your fabric using Fabric Manager Web Services, follow these steps:

Step 1 Choose **Performance > Traffic Analyzer**.

You see a summary table of all SPAN destination ports and configured Traffic Analyzers in your fabric shown in Figure 9-4.

Figure 9-4 Traffic Analyzer in Fabric Manager Web Services

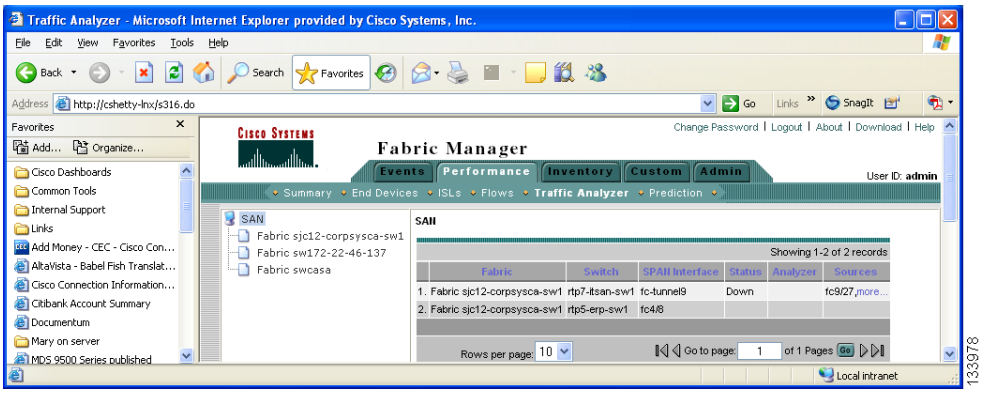

Step 2 Navigate to the fabric or VSAN where you want to rediscover instances of Traffic Analyzer from the navigation bar.

Step 3 Set Search on Subnet to the subnet that you want to rediscover.

Step 4 Click **Rediscover** to find instances of Traffic Analyzer within the selected fabric or VSAN and subnet

Cisco Systems

Send documentation comments to mdsfeedback-doc@cisco.com

PART 2

Cisco MDS SAN-OS Installation and Switch Management

CHAPTER 10

Obtaining and Installing Licenses

Licenses are available in all switches in the Cisco MDS 9000 Family. Licensing allows you to access specified premium features on the switch after you install the appropriate license for that feature.

This chapter contains information related to licensing types, options, procedures, installation, and management for the Cisco MDS SAN-OS software.

This chapter includes the following sections:

Licensing Terminology

The following terms are used in this chapter:

- Licensed feature—Permission to use a particular feature through a license file, a hardware object, or a legal contract. This permission is limited to the number of users, number of instances, time span, and the implemented switch.
- Licensed application—A software feature that requires a license to be used.
- License enforcement—A mechanism that prevents a feature from being used without first obtaining a license.
- Node-locked license—A license that can only be used on a particular switch using the switch's unique host ID.
- Host IDs—A unique chassis serial number that is specific to each Cisco MDS switch.
- Proof of purchase—A document entitling its rightful owner to use licensed feature(s) on one Cisco MDS switch as described in that document. Also known as the claim certificate.
- Product Authorization Key (PAK)—The PAK allows you to obtain a license key from one of the sites listed in the proof of purchase document. After registering at the specified website, you will receive your license key file and installation instructions through e-mail.
- License key file—A switch-specific unique file that specifies the licensed features. Each file contains digital signatures to prevent tampering and modification. License keys are required to use a licensed feature. License keys are enforced within a specified time span.
- Counted license—The number of licenses issued for a single feature (for example, FCIP). You can increase counted licenses (incremental licenses) should a need arise in the future.
- Missing license—If the bootflash has been corrupted or a supervisor module replaced after a license has been installed, that license will show as "missing." The feature will still work, but the license count will be inaccurate. You should reinstall the license as soon as possible.
- Incremental license—An additional licensed feature that was not in the initial license file. License keys are incremental—if you purchase some features now and others later, the license file and the software detect the sum of all features for the specified switch.
- Evaluation license—A temporary license. Evaluation licenses are time bound (valid for a specified number of days) and are not tied to a host ID (switch serial number).
- Grace period—The amount of time the features in a license package can continue functioning without a license.
- Support—If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

Licensing Model

The licensing model defined for the Cisco MDS product line has two options:

- Feature-based licenses allow features that are applicable to the entire switch. Table 10-1 lists the feature-based license packages.
- Module-based licenses allow features that require additional hardware modules. An example is the IPS-8 or IPS-4 module using the FCIP feature.

Note Each module requires its own separate license. If you replace a module that requires a license with a module of the same type (such as replacing a Storage Services Module (SSM) with another SSM), the existing license will support the new module.

Note Any feature not included in a license package is bundled with the Cisco MDS 9000 Family and is provided at no extra charge to you.

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

Table 10-1 Feature-Based Licenses

Feature License	Features
Enterprise package (ENTERPRISE_PKG)	• Enhanced security features: – LUN zoning – Read-only zones • Port security • VSAN-based access control • Fibre Channel Security Protocol (FC-SP) authentication • Advanced traffic engineering—quality of service (QoS) • IP Security Protocol (IPsec) for iSCSI and FCIP using the MPS-14/2 module or Cisco MDS 9216i switch. • IKE digital certificates • Extended credits using the MPS-14/2 module or the Cisco MDS 9216i Switch • Enhanced VSAN routing—inter-VSAN routing (IVR) over FC • IVR Network Address Translation (NAT) over FC **Note** The FCIP and IVR features are bundled with the Cisco MDS 9216i Switch and do not require the Enterprise package. • Zone-based traffic prioritizing • Zone-based QoS • Extended credits • Fibre Channel write acceleration • SCSI flow statistics • FCIP Encryption • Fabric binding for Fibre Channel • FCIP encryption

Table 10-1 Feature-Based Licenses (continued)

Feature License	Features
SAN extension over IP package for IPS-8 modules (SAN_EXTN_OVER_IP) SAN extension over IP package for IPS-4 modules (SAN_EXTN_OVER_IP_IPS4)	The following features apply to IPS-8 and IPS-4 modules: • FCIP • FCIP compression • FCIP write acceleration • FCIP tape acceleration • SAN extension tuner features • IVR over FCIP • IVR NAT over FCIP
SAN extension over IP package for MPS-14/2 modules (SAN_EXTN_OVER_IPS2) **Note** The FCIP, IVR, and SAN extension tuner features are bundled with the Cisco MDS 9216i switch and do not require the SAN extension over IP package to be installed for the fixed IP ports on the integrated supervisor module.	The following features apply to the MPS-14/2 module: • FCIP • FCIP compression • FCIP write acceleration • FCIP tape acceleration • SAN extension tuner features • IVR over FCIP • IVR NAT over FCIP • Hardware-based FCIP compression
Mainframe package (MAINFRAME_PKG)	• FICON protocol and CUP management • FICON VSAN and intermixing • FICON tape acceleration • FICON port numbering • Switch cascading • Fabric binding • IBM TotalStorage Virtual Tape Server (VTS) • IBM TotalStorage XRC application

Table 10-1 Feature-Based Licenses (continued)

Feature License	Features
Fabric Manager Server package (FM_SERVER_PKG)	• Multiple physical fabric management • Centralized fabric discovery services • Continuous MDS health and event monitoring • Long term historical Fibre Channel performance monitoring • Custom performance reports and charting for hotspot analysis • Performance predition • Performance threshold monitoring • Fabric Manager Web Client for operational view • Fabric Manager server proxy services • Server performance summary report • Configurable RRD collection parameters • Data collection auto update • Event forwarding • Filtering by user-defined groups
Storage Services Enabler package (STORAGE_SERVICES_ENABLER_PKG)	• The underlying infrastructure and programmatic interface to enable network-hosted storage applications when used with the Storage Services Modules (SSMs). • The intelligent fabric applications running on the SSM that require the SSE license are as follows: – SANTap – Network-Accelerated Serverless Backup (NASB) – Third-party partner applications

Licensing High Availability

As with other Cisco MDS SAN-OS features, the licensing feature also maintains the following high availability standards for all switches in the Cisco MDS 9000 Family:

- Installing any license in any switch is a nondisruptive process.
- Installing a license automatically saves a copy of permanent licenses to the chassis in all switches.
- Enabling a license feature without a license key starts a counter on the grace period. You then have 120 days to install the appropriate license keys or disable the use of that feature. If at the end of the 120-day grace period the switch does not have a valid license key for the feature, the feature is automatically disabled by the switch.

Directors in the Cisco MDS 9500 Series have the following additional high availability features:

- The license software runs on both supervisor modules and provides failover protection.

- The license key file is mirrored on both supervisor modules. Even if both supervisor modules fail, the license file continues to function from the version that is available on the chassis.

Options to Install a License

If you have purchased a new switch through either your reseller or through Cisco Systems, you can:

- Obtain a factory-installed license (only applies to new switch orders).
- Perform a manual license installation (applies to existing switches).

Obtaining a Factory-Installed License

You can obtain factory-installed licenses for a new switch.

To obtain a factory-installed license for a new Cisco MDS switch, follow these steps:

Step 1 Contact your reseller or Cisco representative and request this service.

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

Your switch is shipped with the required licenses installed in the system. The proof of purchase document is sent along with the switch.

Step 2 Obtain the host ID from the proof of purchase document for future use.

Step 3 Start to use the switch and the licensed features.

Performing a Manual Installation

If you have existing switches or if you wish to install the licenses on your own, you must first obtain the license key file and then install that file in the switch (see Figure 10-1).

Figure 10-1 Obtaining a License Key File

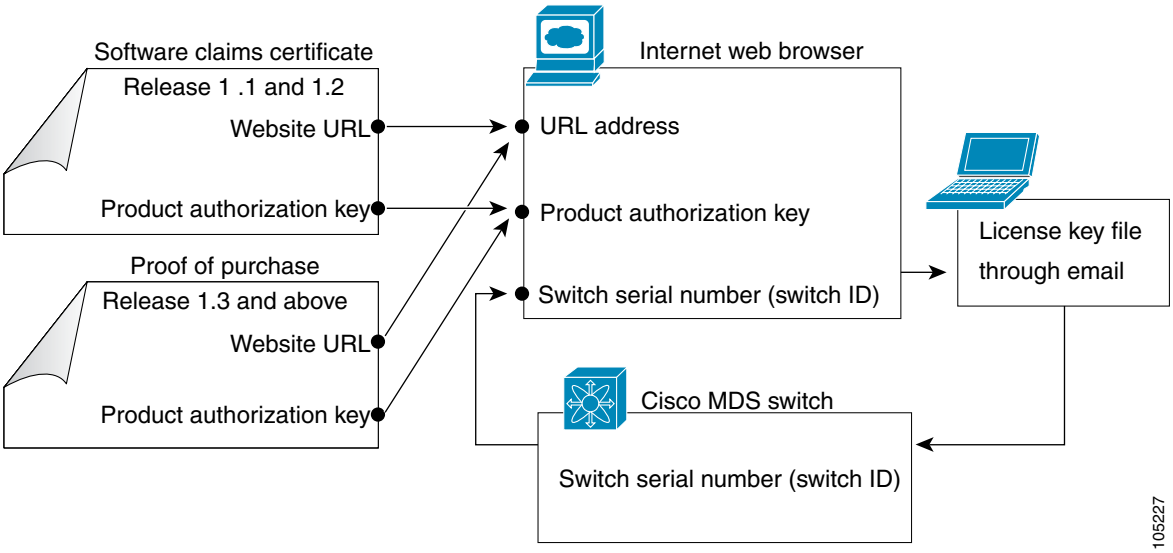

Obtaining the License Key File

To obtain new or updated license key files using Device Manager, follow these steps:

Step 1 Select **Physical > Inventory** from the main menu. You see the inventory for the switch. The host ID is referred to as the serial number.

Tip Prepend the serial number with VDH=. For example, if the serial number is `FOX064317SQ`, the full serial number is `VDH=FOX064317SQ`.

Step 2 Obtain either your claim certificate or your proof of purchase document. This document accompanies every Cisco MDS switch.

Step 3 Get the product authorization key (PAK) from either the claim certificate or the proof of purchase document.

Step 4 Locate the website URL from either the claim certificate or the proof of purchase document.

Step 5 Access the specified URL that applies to your switch and enter the switch serial number and the PAK.

The license key file is sent to you by e-mail. The license key file is digitally signed to only authorize use on the requested switch. The requested features are also enabled once the Cisco SAN-OS software on the specified switch accesses the license key file.

Caution Install the license key file in the specified MDS switch without making any modifications.

A license is either permanent or it expires on a fixed date. If you do not have a license, the grace period for using that feature starts from the first time you start using a feature offered by that license (see the "Grace Period Alerts" section on page 10-14).

Step 6 Use the **copy licenses** CLI command in EXEC mode to save your license file to one of two locations—the bootflash: directory or the slot0: device. Refer to the *Cisco MDS 9000 Family CLI Configuration Guide* for more information.

Installing the License Key File

Tip If you need to install multiple licenses in any switch in the Cisco MDS 9000 Family, be sure to provide unique file names for each license key file.

. The best way to install licenses on the switches in your fabric is to use the License Wizard provided in Fabric Manager. You can also use Device Manager to install licenses on each switch individually.

Note You do not need a license to access a switch with Fabric Manager. See the "Licensing Model" section on page 10-3 for a list of features requiring licenses.

You can install licenses two ways:

- Using the Fabric Manager License Wizard
- Using Device Manager

Installing Licenses Using Fabric Manager License Wizard

To install licenses using the Fabric Manager License Wizard, follow these steps:

Step 1 Log into a switch in the fabric containing the switches for which you want to install licenses.

To install licenses on multiple switches, you do not need to log into each switch; however, the switches must be in the fabric you are viewing.

Step 2 Start the License Wizard by clicking **Tools > Other >License Install**. Or, you can select **Licenses** under **Switches** in the Physical Attributes pane.

You see the license information in the Information pane, one line per feature.

Step 3 Click the **Keys** tab, and then click the **License Install Wizard** icon in the toolbar.

Figure 10-2 License Install Wizard Icon

You see the initial screen of the License Wizard.

Step 4 If you have already obtained the license key files, click the corresponding radio button and proceed to Step 6.

Step 5 Click **I have the Product Authorization Key (PAK)** if you have the authorization key.

Step 6 Select the vendor, from whom you purchased your switch, in the Vendor drop-down list.

The License Server URL changes depending on the vendor you select. If your URL is different, or if you select **Other** as the vendor, enter the correct license server URL.

Note In some cases, license validation from Cisco partners requires Java version 1.4.2_04 or higher. If you cannot install licenses from a Cisco partner, check to make sure your Java version is at least 1.4.2_04.

Step 7 Click **Next** to continue to the next screen.

Step 8 Select the switches for which you have PAKs or license key files.

When you check the check box for a switch, the PAK or license file name field for that switch becomes editable. The VDH=<serial number> for each switch is shown in the Host ID column.

Step 9 Enter the PAK or license file name for each switch you have selected in the appropriate column. If you have the license files on your PC, you can double-click in the License File Name text area to bring up a dialog box and browse for the license files.

You can install multiple licenses on the same switch using different PAKs. To do this, enter the PAKs separated by commas.

Step 10 Click **Finish** to transfer the licenses from the host to the switches.

Fabric Manager accesses the appropriate license site and installs the licenses onto each switch. The status of each installation is displayed in the Status column, as follows:

- success—Install or uninstall operation completed successfully.
- inProgress—License install or uninstall operation is in progress.
- corruptedLicenseFile—License file content is invalid or corrupted.
- targetLicenseFileAlreadyExist—Target license file name already exists.
- invalidLicenseFileName—License file does not exist.
- duplicateLicense—License file is already installed.
- generalLicensingFailure—General error from License Manager.
- none—No install operation is performed.
- licenseExpiryConflict—License exists with a different expiration date for the feature.
- invalidLicenseCount—License count is invalid for the feature.

Step 11 Click the **Close** button to close the wizard. To install more licenses at this point, you must close the wizard and launch it again.

Installing or Updating Licenses Using Device Manager

To install a license on your switch using Device Manager, follow these steps:

Step 1 Select **Licenses** from the Admin menu.

You see the Licenses dialog box.

Step 2 Click the **Install** tab.

The HostId shows the "VDH=" portion of the serial number. The rest of the number is completed in Steps 3 through 5.

Step 3 Enter the uniform resource identifier (URI) from which the license file will be retrieved.

You should already have copied the license file provided by Cisco.com or by some other means (for example, through the CLI) to this location.

Step 4 Enter the target file name in the Target Filename field to specify where the license file will be installed.

Step 5 Click **Install** if you are installing, or **Update** if you are updating.

You see the status of the installation at the bottom of the dialog box, as follows:

- success—Install or uninstall operation completed successfully.
- inProgress—License install or uninstall operation is in progress.
- corruptedLicenseFile—License file content is invalid or corrupted.
- targetLicenseFileAlreadyExist—Target license file name already exists.
- invalidLicenseFileName—License file does not exist.

- duplicateLicense—License file is already installed.
- generalLicensingFailure—General error from License Manager.
- none—No install operation is performed.
- licenseExpiryConflict—License exists with a different expiration date for the feature.
- invalidLicenseCount—License count is invalid for the feature.
- notThisHost—License host ID in the license file doesn't match.
- licenseInGraceMore—Number of licenses in grace period is more than the number in the install license file.
- licenseFileNotFound—License file not found for the install, uninstall, or update operation.
- licenseFileMissing—A previously installed license file is found missing.
- invalidLicenseFileExtension—License file does not have a .lic extension.
- invalidURI—Invalid license file URI specified for install operation.
- noDemoLicenseSupport—Demo license not supported.
- invalidPlatform—Invalid platform.

Step 6 Repeat Steps 3 through 5 to install another license, or click **Close** to close the License Manager dialog box.

Identifying License Features in Use

When a feature is enabled, it can activate a license grace period.

To identify the features active for a specific license using Fabric Manager, follow these steps:

Step 1 Select a switch from the Fabric pane, or select a group of switches (SAN, fabric, VSAN) from the Logical Domains pane.

Step 2 Select **Licenses** under **Switches in** the Physical Attributes pane.

You see the contents of the **Feature Usage** tab in the Information pane, with installed licenses listed in the Feature column.

Step 3 Click the **Usage** tab.

You see the features currently in use in the Application column.

Uninstalling Licenses

You can only uninstall a permanent license that is not in use. If you try to delete a permanent license that is currently being used, the software rejects the request and issues an error message. Uninstalling an unused license causes the grace period to come into effect. The grace period is counted from the first use of the feature without a license and is reset when a valid license file is installed.

Note Permanent licenses cannot be uninstalled if they are currently being used. Features turned on by permanent licenses must first be disabled, before that license is uninstalled.

Tip If you are using an evaluation license and would like to install a new permanent license, you can do so without service disruption and before the evaluation license expires. Removing an evaluation license immediately triggers a grace period without service disruption.

Caution Disable related features before uninstalling a license. The delete procedure fails if the license is in use.

To uninstall a license, follow these steps:

Step 1 Log into the switch. If you are using Fabric Manager to remove licenses from multiple switches, you do not need to log in to each switch; however, the switches must be in the fabric you are viewing.

Step 2 From the Fabric Manager Physical Attributes pane, select **Licenses** under **Switches**. You see the license information in the Information pane, one line per feature.

From Device Manager, click **Admin > Licenses** from the menu. You see the Licenses dialog box.

Step 3 In Fabric Manager, click the **Keys** tab. You see the list of License Key files. Click the name of the license you want to remove, and press the Delete keyboard key or click the **Delete Row** icon in the toolbar.

In Device Manager, click **Uninstall**, and enter the name of the License Key file you want to remove. Click **Apply** to remove the License Key file, and click **Close** to close the dialog box.

Note To delete a license, you must disable the features enabled by that license. The delete procedure fails if the license is in use, and an error message is displayed.

Updating Licenses

If your license is time bound, you must obtain and install an updated license. Contact technical support to request an updated license.

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

To update a license, follow these steps:

Step 1 Obtain the updated license file using the procedure described in the "Obtaining the License Key File" section on page 10-9.

Step 2 Save your running configuration to a remote server using the **copy** command (see the "Copying Files" section on page 14-5).

Step 3 Verify the name of the file to be updated.

Step 4 Follow the procedure for updating a license described in the "Installing or Updating Licenses Using Device Manager" section on page 10-11.

Grace Period Alerts

Cisco SAN-OS gives you a 120 day grace period. This grace period starts or continues when you are evaluating a feature for which you have not installed a license.

The grace period stops if you disable a feature you are evaluating, but if you enable that feature again without a valid license, the grace period countdown continues where it left off.

The grace period operates across all features in a license package. License packages can contain several features. If you disable a feature during the grace period and there are other features in that license package that are still enabled, the countdown does not stop for that license package. To suspend the grace period countdown for a license package, you must disable every feature in that license package.

The Cisco SAN-OS license counter keeps track of all licenses on a switch. If you are evaluating a f feature and the grace period has started, you will receive console messages, SNMP traps, system messages, and Call Home messages on a daily basis.

Beyond that, the frequency of these messages become hourly during the last seven days of the grace period. The following example uses the FICON feature. On January 30th, you enabled the FICON feature, using the 120 day grace period. You will receive grace period ending messages as:

- Daily alerts from January 30th to May 21st.
- Hourly alerts from May 22nd to May 30th.

On May 31st, the grace period ends, and the FICON feature is automatically disabled. You will not be allowed to use FICON until you purchase a valid license.

Note You cannot modify the frequency of the grace period messages.

Caution After the final seven days of the grace period, the feature is turned off and your network traffic may be disrupted. Any future upgrade will enforce license requirements and the 120-day grace period.

License Transfers Between Switches

A license is specific to the switch for which it is issued and is not valid on any other switch. If you need to transfer a license from one switch to another, contact your customer service representative.

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

Displaying License Information

Use Fabric Manager or Device Manager to display all license information configured on this switch.

Viewing License Information in Fabric Manager

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

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

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

Step 3 Click the **Keys** tab to display the information about each of the License Key files installed on your switches.

Caution Once an expiration period has started, notifications appear in the Fabric Manager's Events pane on a daily basis. During the last seven days of the expiration period, these messages are displayed hourly. After the final seven days of the expiration period, the feature is turned off and your network traffic may be disrupted.

Step 4 Click the **Usage** tab to see the applications using the feature package on each switch. Use this tab to determine which applications depend on each license installed.

Viewing License Information in Device Manager

To view license information in Device Manager, follow these steps:

Step 1 Click **Admin > Licenses** from the menu.

You see the Licenses dialog box.

Step 2 Click the **Features** tab to see the name of the feature package, the type of license, the expiration date, the grace period (if you do not have a license for a particular feature), and any errors, such as a missing license.

Step 3 Click the **Files** tab to display the information about each of the License Key files installed on your switch.

Step 4 Click the **Install** tab to install or update a license file.

Step 5 Click the **Usage** tab to which applications are using the features on the switch.

Viewing Licenses Using Fabric Manager Web Services

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

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

Fabric Manager Server Licensing

When you install Fabric Manager, the a basic version of the Fabric Manager Server (FMServer) is installed with it. To get the enhanced features, such as Performance Manager and remote client support) you will need to buy and install the Cisco MDS 9000 Family Fabric Manager Server license package.

However, trial use of these enhanced features is available. To enable the 120-day trial, you simply use the feature as you would if you had purchased the license. You see a dialog box explaining that this is a demo version, enabled for a limited time.

If you are evaluating Fabric Manager Server features and want to stop the evaluation period for that feature, you can do that using Device Manager.

To stop the evaluation using Device Manager, follow these steps:

Step 1 Click **Admin > Licenses**.

You see the Licenses dialog box.

Step 2 Click the **Features** tab and select the feature to check in.

When you select the feature, you see a Check In FM button at the bottom of the dialog box.

Step 3 Click **Check In FM** to stop the demo period timer.

CHAPTER 11

Initial Configuration

This chapter describes how to set up some of the initial configuration parameters for switches using Fabric Manager. Most of the initial switch configuration procedures can only be performed using the CLI. Refer to the *Cisco MDS 9000 Family CLI Configuration Guide* for this information. This chapter includes the following sections:

Assigning a Switch Name

Each switch in the fabric requires a unique name. You can assign names to easily identify the switch by its physical location, its SAN association, or the organization to which it is deployed. The assigned name is displayed in the command-line prompt. The switch name is limited to 20 alphanumeric characters.

To change the name of a switch using Fabric Manager, follow these steps:

Step 1 Expand **SAN** in the Logical Domains pane, select a fabric or a VSAN from the Logical Domains pane.

Step 2 Expand **Switches** in the Physical Attributes pane.

You see a list of switches in the Information pane.

Step 3 Double-click the Logical Name of the switch you want to change in the Information pane.

You see the name highlighted with a blinking cursor next to it (see Figure 11-1).

Figure 11-1 Changing the Logical Name of a Switch

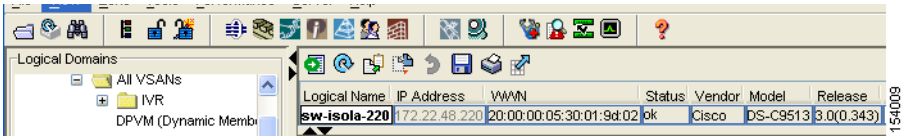

Step 4 Type the new name of the switch (see Figure 11-1).

Step 5 Click the **Apply Changes** icon.

Step 6 Right-click the Fabric pane map and choose **Refresh** to see your changes.

Verifying the Module Status

Before you begin configuring the switch, you need to ensure that the modules in the chassis are functioning as designed.

To verify the status of a module at any time, follow these steps:

Step 1 Expand **SAN** in the Logical Domains pane, then select a fabric or a VSAN from the Logical Domains pane.

Step 2 Expand **Switches** and choose **Hardware** in the Physical Attributes pane.

You see the contents of the **Inventory** tab in the Information pane shown in Figure 11-2.

Figure 11-2 Inventory of a Selected Module

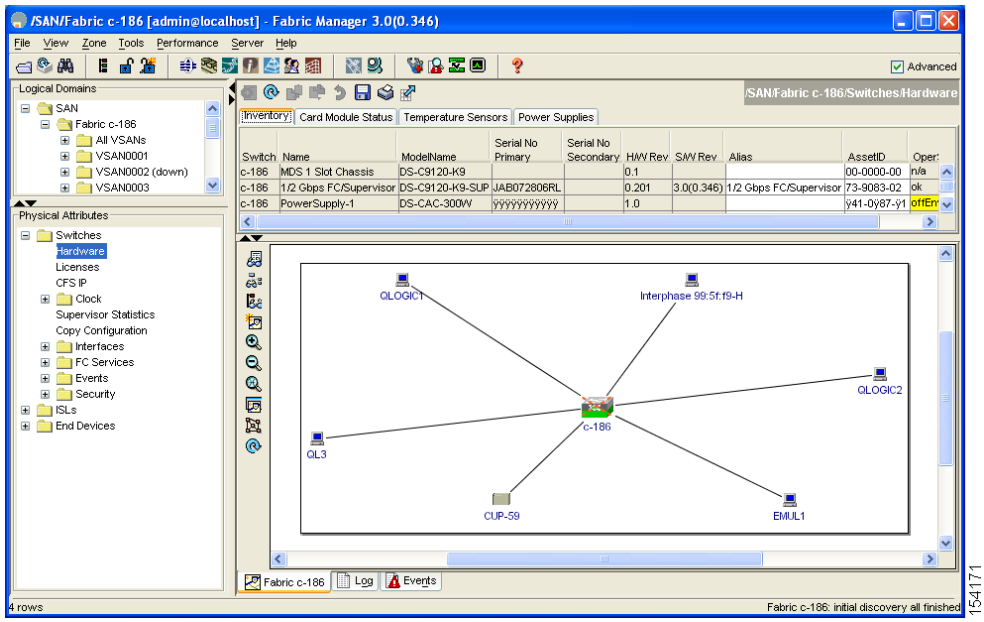

Step 3 Click the **Card Module Status** tab.

You see the status in the Oper Status column of each module in each switch of the SAN, fabric, or VSAN you selected.

Figure 11-3 Card Module Status

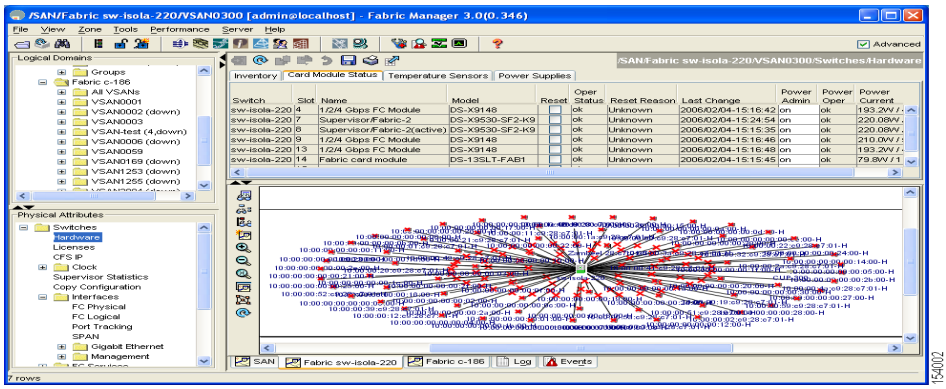

If the status is OK or active, continue with your configuration (see Chapter 17, "Managing Modules").

Configuring Date, Time, and Time Zone

Switches in the Cisco MDS 9000 Family use Universal Coordinated Time (UTC), which is the same as Greenwich Mean Time (GMT).

To change the default time on the switch with Fabric Manager, follow these steps:

Step 1 Expand **SAN**, then select a fabric or a VSAN in the Logical Domains pane.

You see a list of switches in the Information pane.

Step 2 Expand **Switches** and select **Clock** in the Physical Attributes pane.

You see the clock information in the Information pane shown in Figure 11-4.

Figure 11-4 Clock Date and Time for Selected Switch

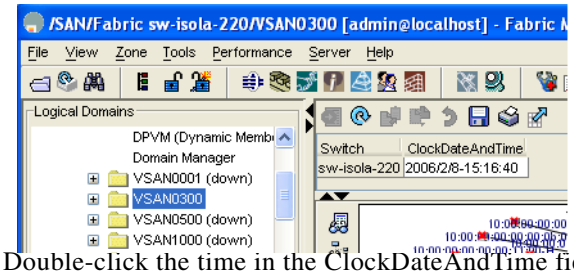

Step 3 Double-click the time in the ClockDateAndTime field for the switch to change.

Step 4 Enter the date, time, and time zone in the format *YYYY/MM/DD-hh:mm:ss* ZONE,

Where:

- *YYYY* is the year (2002)
- *MM* is the month (08)
- *DD* is the date (23)
- *hh* represents hours in military format (15 for 3 p.m.)
- *mm* is minutes (58)

- *ss* is seconds (09)
- ZONE is the time zone (GMT)

Note If you do not enter a time zone, GMT is used as the default.

Step 5 Click the **Apply Changes** icon.

Note The date and time changes are saved across system resets.

NTP Configuration

A Network Time Protocol (NTP) server provides a precise time source (radio clock or atomic clock) to synchronize the system clocks of network devices. NTP is transported over User Datagram Protocol UDP/IP. All NTP communications use Universal Time Coordinated (UTC). An NTP server receives its time from a reference time source, such as a radio clock or atomic clock, attached to the time. NTP distributes this time across the network.

In a large enterprise network, having one time standard for all network devices is critical for management reporting and event logging functions when trying to correlate interacting events logged across multiple devices. Many enterprise customers with extremely mission-critical networks maintain their own stratum-1 NTP source.

Time synchronization happens when several frames are exchanged between clients and servers. The switches in client mode know the address of one or more NTP servers. The servers act as the time source and receive client synchronization requests.

By configuring an IP address as a peer, the switch will obtain and provide time as required. The peer is capable of providing time on its own and is capable of having a server configured. If both these instances point to different time servers, your NTP service is more reliable. Thus, even if the active server link is lost, you can still maintain the right time due to the presence of the peer.

Tip If an active server fails, a configured peer helps in providing the NTP time. Provide a direct NTP server association and configure a peer to ensure backup support if the active server fails.

If you only configure a peer, the most accurate peer takes on the role of the NTP server and the other peer(s) acts as a peer(s). Both machines end at the right time if they have the right time source or if they point to the right NTP source.

You can configure NTP using either IPv4 addresses, IPv6 addresses, or DNS names.

Create an NTP Server or Peer

To create an NTP server or peer, follow these steps:

Step 1 In the Fabric Manager Physical pane, expand **Switches** then select **System**, or from Device Manager, choose **Admin > NTP**.

In Fabric Manager, you see the System information pane. In Device Manager, you see the NTP dialog box.

Figure 11-5 Device Manager NTP Dialog Box

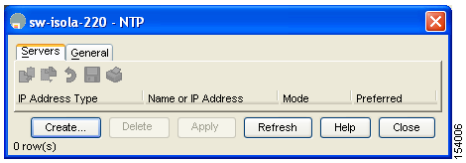

Step 2 Click the **NTP Peer** tab.

You see a list of NTP peers and servers for that switch.

Step 3 Click **Create**.

You see the Create NTP Peer dialog box.

Step 4 Enter the peer address in the Peer Address field.

Step 5 Choose the mode (**peer** or **server**).

Step 6 Check the **Preferred** check box if you want this peer to be a Preferred Peer.

Step 7 Click **Create** to create the peer or server, or click **Close** to close the dialog box without creating the peer or server.

The new peer or server is listed on the Peer tab.

Edit an NTP Server or Peer Configuration

To edit an NTP server or peer, follow these steps.

Step 1 In the Fabric Manager Physical Attributes pane, expand **Switches** then select **System**, or from Device Manager, choose **Admin > NTP**.

In Fabric Manager, you see the System information pane. In Device Manager, you see the NTP dialog box.

Step 2 Click the **NTP Peer** tab.

You see a list of NTP peers and servers for that switch.

Step 3 Change the peer address by double-clicking the IP address in the Peer Address column, and changing the numbers. Alternatively, you can triple click the IP address and type in a new address.

Step 4 Change the switch mode from **peer** to **server** by clicking the Mode column next to the address of the switch.

You see a drop-down list. Select the mode (**peer** or **server**) you want for the switch.

Step 5 Change the peer status of the switch to Preferred Peer by checking the **PrefPeer** check box next to the address of the switch. To remove this status, uncheck the check box.

Step 6 Click **Apply** to apply your changes to the switch, or click **Close** to close the dialog box without saving your changes.

Delete an NTP Server or Peer

To delete an NTP server or peer, follow these steps.

Step 1 In the Fabric Manager Physical pane, expand **Switches** and choose **System**, or from Device Manager, choose **Admin > NTP**.

In Fabric Manager, you see the System information pane. In Device Manager, you see the NTP dialog box.

Step 2 Click the **NTP Peer** tab.

You see a list of NTP peers and servers for that switch.

Step 3 Delete a server or peer by clicking the IP address in the Peer Address column.

The Delete button is enabled.

Step 4 Click **Delete** to delete the peer or server, or click **Close** to close the dialog box without deleting the peer.

NTP Configuration Guidelines

The following guidelines apply to all NTP configurations:

- You should have a peer association with another switch only when you are sure that your clock is reliable (which means that you are a client of a reliable NTP server).
- A peer configured alone takes on the role of a server and should be used as backup. If you have two servers, then you can have several switches point to one server, and the remaining switches to the other server. Then you would configure peer association between these two sets. This forces the clock to be more reliable.
- If you only have one server, it's better for all the switches to have a client association with that server.

Not even a server down time will affect well-configured switches in the network. Figure 11-6 displays a network with two NTP stratum 2 servers and two switches.

Figure 11-6 NTP Peer and Server Association

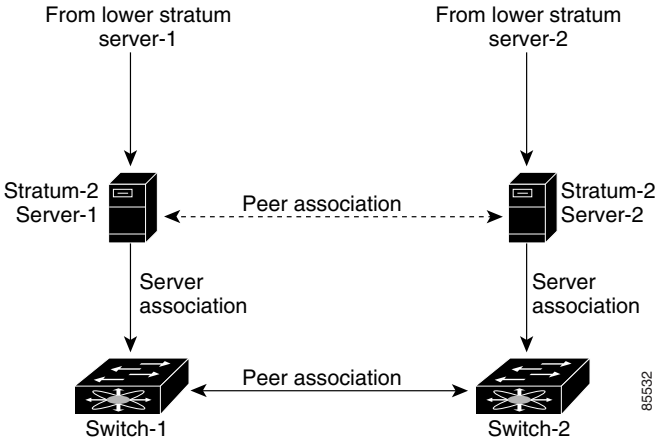

In this configuration, the switches were configured as follows:

- Stratum 2 Server 1
  - IP address -10.10.10.10
  - Stratum-2 Server-2
  - IP address -10.10.10.9
- Switch 1 IPv4 address–10.10.10.1
- Switch 1 NTP configuration
  - NTP server 10.10.10.10
  - NTP peer 10.10.10.2
- Switch 2 IP address -10.10.10.2
- Switch 2 NTP configuration
  - NTP server 10.10.10.9
  - NTP peer 10.10.10.1

NTP Configuration Distribution

You can enable NTP fabric distribution for all Cisco MDS switches in the fabric. When you perform NTP configurations, and distribution is enabled, the entire server/peer configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you issue the first configuration command after you enabled distribution in a switch. The NTP application uses the effective and pending database model to store or commit the commands based on your configuration.

See Chapter 12, “Using the CFS Infrastructure” for more information on the CFS application.

Committing NTP Configuration Changes

When you commit the NTP configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. When you commit the NTP configuration changes without implementing the session feature, the NTP configurations are distributed to all the switches in the fabric.

Discarding Configuration Changes

After making the configuration changes, you can choose to discard the changes or to commit them. In either case, the lock is released.

Configure NTP with CFS

To configure NTP with CFS using Fabric Manager, follow these steps:

Step 1 Expand **Switches,** expand **Clock** then select **NTP** in the Physical Attributes pane.

You see the feature configuration in the Information pane.

Step 2 Click the **CFS** tab in the Information pane.

You see the CFS configuration and status for each switch.

Step 3 Click a switch value in the Global column, **enable** or **disable**.

A drop-down menu appears (see Figure 11-7).

Figure 11-7 Enabling or Disabling NTP with CFS for a Switch

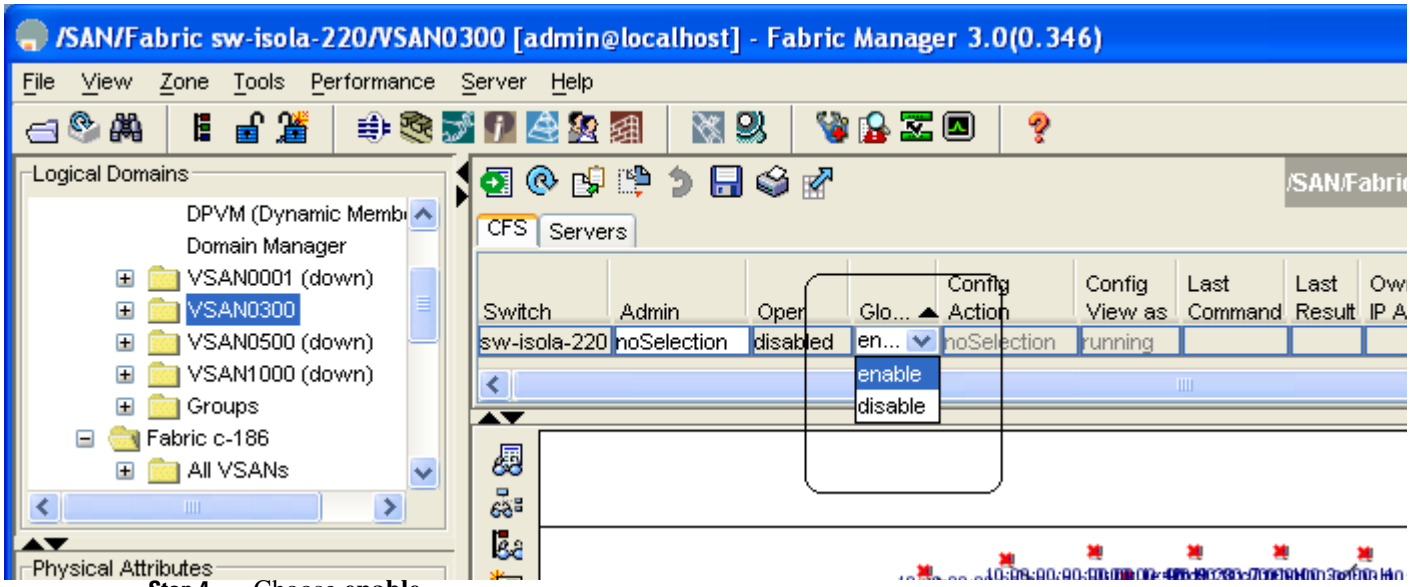

Step 4 Choose **enable.**

Step 5 Repeat steps 3 and 4 for all switches in the fabric.

Note A warning displays if you do not enable CFS for all switches in the fabric for this feature.

Step 6 Check the **Master** check box for the switch that you want to act as the merge master for this feature.

Step 7 Click the switch value in the Config Action column. A drop-down menu appears.

Step 8 Select **Commit**.

Step 9 Click the **Servers tab in** the Information pane. You see the configuration for this feature based on the master switch.

Step 10 Modify the Master configuration as needed. For example, right-click the value in the Master column and select **Create Row** to create a server for NTP.

a. Set the ID, and the Name or IP Address for the NTP server.

b. Choose a **Mode** radio button and, optionally, check the **Preferred** check box.

c. Click **Create** to add the server.
Fabric Manager sends the request to the master switch. Click the **CFS** tab and check the Last Results column for the new entry. It has a "pending" status.

Step 11 From the **CFS** tab, set the Config Action column to **commit** to distribute the feature change through the fabric. Fabric Manager only changes the status to “running” when **commit**, **clear,** or **abort** is selected and applied.

Note Fabric Manager will not change the status to “pending” if **enable** is selected, because the “pending” status does not apply until the first actual change is made.

Step 12 Click the **Apply Changes** icon to commit the configuration changes for that feature and distribute the changes through CFS, or click **Undo Changes** to discard the changes for that feature.

Releasing Fabric Session Lock

If you have performed an NTP fabric task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.

Tip The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

Database Merge Guidelines

When merging two fabrics, follow these guidelines:

- Be aware that the merge is a union of the existing and the received database in each switch in the fabric.
- Do not configure an IP address as a server on one switch and as a peer on another switch. The merge can fail if this configuration exists.
- Verify that the union of the databases does not exceed the maximum limit of 64.

See the "CFS Merge Support" section on page 12-9 for detailed concepts.

Management Interface Configuration

The management interface on the switch allows multiple simultaneous Telnet or SNMP sessions. You can remotely configure the switch through the management interface (mgmt0), but first you must configure some IP parameters so that the switch is reachable. You can manually configure the management interface from the CLI. You can configure the mgmt 0 interface with either IPv4 address parameters or an IPv6 address.

On director class switches, a single IP address is used to manage the switch. The active supervisor module's mgmt0 interface uses this IP address. The mgmt0 interface on the standby supervisor module remains in an inactive state and cannot be accessed until a switchover happens. After a switchover, the mgmt0 interface on the standby supervisor module becomes active and assumes the same IP address as the previously active supervisor module.

The mgmt0 is autosensing and operates in full duplex mode at a speed of 10/100/1000 Mbps (1000 Mbps is only available on the Supervisor-2 module). The speed and mode cannot be configured. Autosensing supports both the speed and the duplex mode.

You can set the management interface in the Fabric Manager Preferences screen to use SNMP over TCP. The advantages of this setting are an increased buffer size and faster transfer rate. If your fabric has a long timeout period, it may prevent you from using SNMP (which may have a relatively shorter timeout period). If so, change this setting to **false** and restart Fabric Manager Server. UDP is used instead.

Note If it is set to false, the same choice must be set in FabricManager.{shlbat}

Note Before you begin to configure the management interface manually, obtain the switch's IPv4 address and IPv4 subnet mask or the IPv6 address. Make sure the console cable is connected to the console port.

Default Gateway Configuration

The supervisor module sends IP packets with unresolved destination IPv4 addresses to the default gateway (see Figure 11-8).

Figure 11-8 Default Gateway

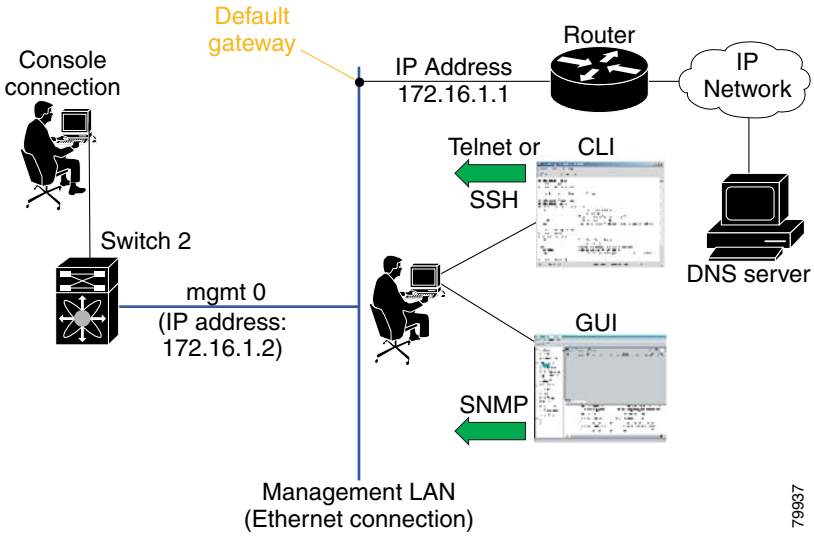

Telnet Server Connection

The Telnet server is enabled by default on all switches in the Cisco MDS 9000 Family. If you require a secure SSH connection, you need to disable the default Telnet connection and then enable the SSH connection (see the "Enabling SSH or Telnet Service" section on page 33-17).

Note For information on connecting a terminal to the supervisor module console port, refer to the *Cisco MDS 9500 Series Hardware Installation Guide* or the *Cisco MDS 9200 Series Hardware Installation Guide*.

Tip A maximum of 16 sessions are allowed in any switch in the Cisco MDS 9500 Series or the Cisco MDS 9200 Series.

Make sure the terminal is connected to the switch and that the switch and terminal are both powered on.

Disabling a Telnet Connection

To disable Telnet connections to the switch using Device Manager, follow these steps:

Step 1 Select **Device** > **Preferences**.

Step 2 Check the **Use Secure Shell instead of Telnet** check box.

Step 3 Click **Apply**.

Telnet is disabled and SSH is enabled on the switch.

Configuring CDP

The Cisco Discovery Protocol (CDP) is an advertisement protocol used by Cisco devices to advertise itself to other Cisco devices in the same network. CDP runs on the data link layer and is independent of Layer 3 protocols. Cisco devices that receive the CDP packets cache the information to make it is accessible through the CLI and SNMP.

CDP is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interfaces on the IPS and MPS-14/2 modules. The CDP daemon is restartable and switchable. The running and startup configurations are available across restarts and switchovers.

CDP version 1 (v1) and version 2 (v2) are supported in Cisco MDS 9000 Family switches. CDP packets with any other version number are silently discarded when received.

When the interface link is established, CDP is enabled by default and three CDP packets are sent at one-second intervals. Following this, the CDP frames are sent at the globally-configured refresh interval.

To globally disable CDP using Fabric Manager, follow these steps:

Step 1 Select a switch in the Logical Domains pane.

Step 2 Expand **Switches**, expand **Interfaces**, expand **Management**, and then select **CDP** in the Physical Attributes pane.

You see the CDP information in the Information pane shown in Figure 11-9.

Figure 11-9 Cisco Discovery Protocol

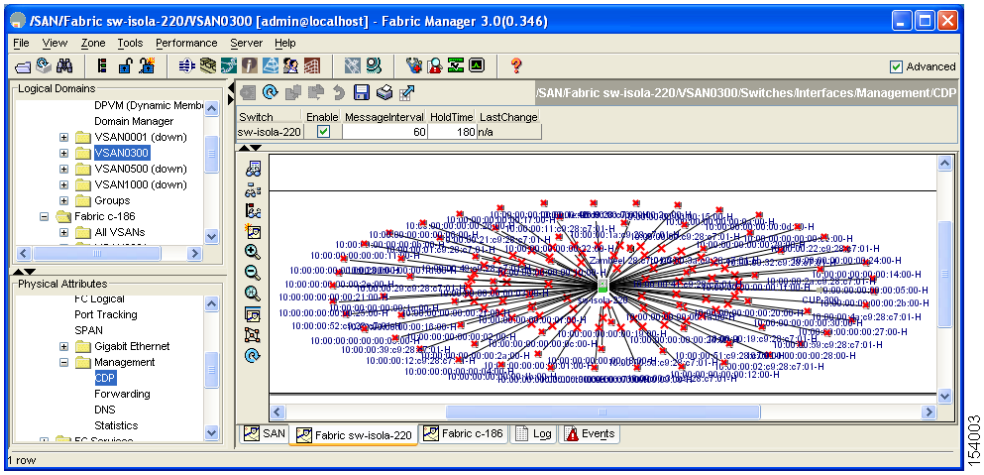

Step 3 Uncheck the **Enable** check box.

Step 4 Click the **Apply Changes** icon.

To disable CDP using Device Manager, follow these steps:

Step 1 Click **IP > CDP**.

You see the CDP dialog box shown in Figure 11-10.

Figure 11-10 Cisco Discovery Protocol

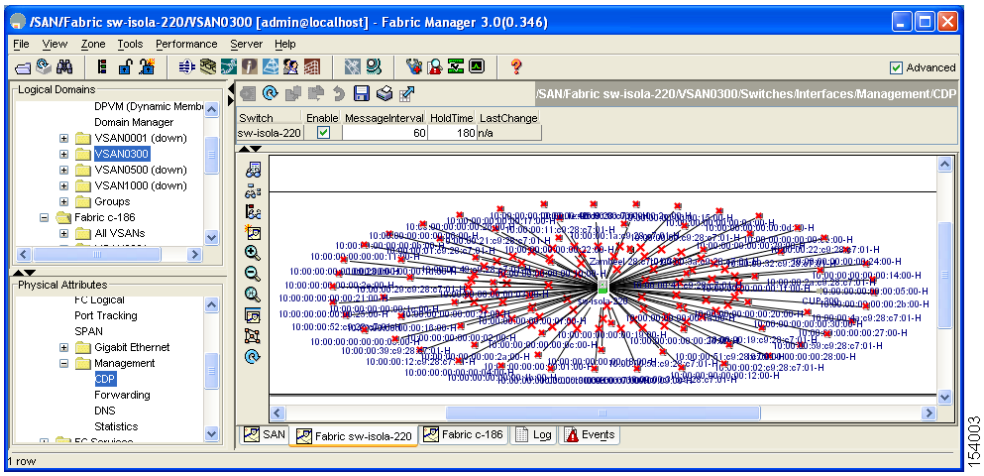

Step 2 Uncheck the **Enable** check box.

Step 3 Click the **Apply Changes** icon.

To globally configure the message interval for the CDP protocol using Device Manager, follow these steps:

Step 1 Click **IP > CDP**.

You see the CDP dialog box shown in Figure 11-11.

Figure 11-11 Cisco Discovery Protocol

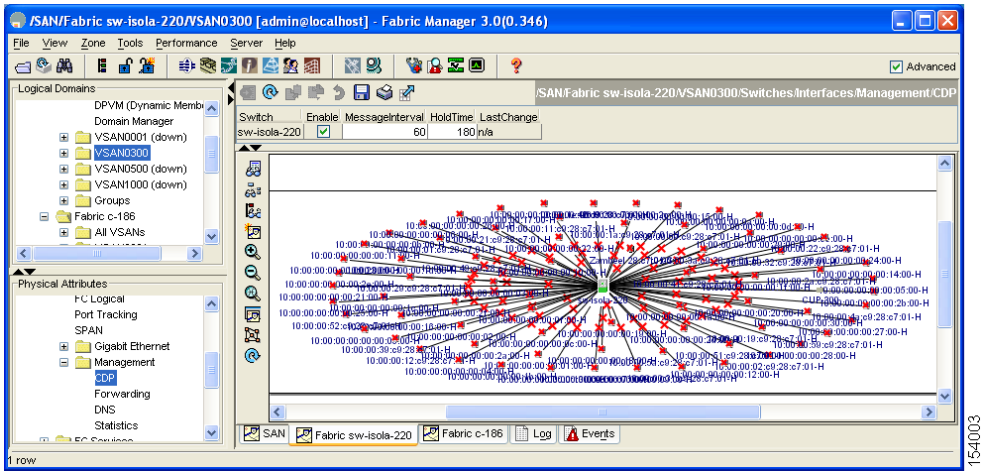

Step 2 Set the message interval time in seconds (5-254).

Step 3 Click the **Apply icon** .

To globally configure the hold time advertised in CDP packets using Device Manager, follow these steps:

Step 1 Click **IP > CDP**.

You see the CDP dialog box shown in Figure 11-12.

Figure 11-12 Cisco Discovery Protocol

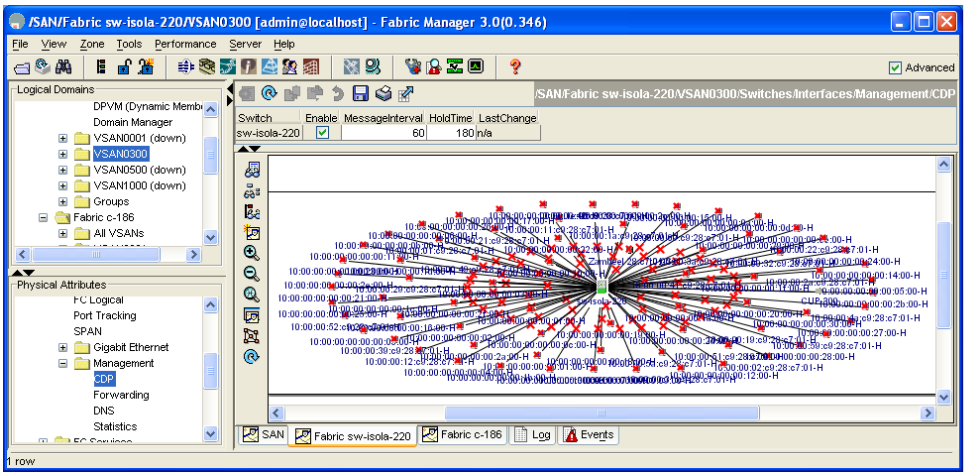

Step 2 Set the hold time in seconds (10-255).

Step 3 Click **Apply**.

CHAPTER 12

Using the CFS Infrastructure

The Cisco MDS SAN-OS software uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database distribution and to foster device flexibility. It simplifies SAN provisioning by automatically distributing configuration information to all switches in a fabric.

Several Cisco MDS SAN-OS applications use the CFS infrastructure to maintain and distribute the contents of a particular application's database.

This chapter contains the following sections:

About CFS

Many features in the Cisco MDS switches require configuration synchronization in all switches in the fabric. Maintaining configuration synchronization across a fabric is important to maintain fabric consistency. In the absence of a common infrastructure, such synchronization is achieved through manual configuration at each switch in the fabric. This process is tedious and error prone.

Cisco Fabric Services (CFS) provides a common infrastructure for automatic configuration synchronization in the fabric. It provides the transport function as well as a rich set of common services to the applications. CFS has the ability to discover CFS capable switches in the fabric and discovering application capabilities in all CFS capable switches.

Cisco SAN-OS Features Using CFS

The following Cisco SAN-OS features use the CFS infrastructure:

- NTP (see the "NTP Configuration Distribution" section on page 11-7).
- Dynamic Port VSAN Membership (see Chapter 24, "Creating Dynamic VSANs").
- Distributed Device Alias Services (see Chapter 27, "Distributing Device Alias Services").
- IVR topology (see "Database Merge Guidelines" section on page 25-32).
- TACACS and RADIUS (see the "Starting a Distribution Session on a Switch" section on page 35-22).
- User and administrator roles (see "Role-Based Authorization" section on page 33-1).
- Port security (see "Port Security Configuration Distribution" section on page 41-17).
- iSNS (see "Configuring iSNS Servers" section on page 45-79).
- Call Home (see "Call Home Configuration Distribution" section on page 58-15).
- Syslog (see "System Message Logging Configuration" section on page 57-3).
- FC timer (see "About fctimer Distribution" section on page 32-12).
- SCSI Flow Services (see "About SCSI Flow Statistics" section on page 49-8).
- Saving startup configurations in the fabric using the Fabric Startup Configuration Manager (FSCM) (see the"Saving Startup Configurations in the Fabric" section on page 14-9).
- Allowed domain ID lists (see the "About Allowed Domain ID Lists" section on page 22-11)
- RSCN timer (see the "Configuring RSCN Timer Distribution Using CFS" section on page 29-7)
- iSLB (see the "About iSLB" section on page 45-42)

CFS Features

CFS has the following features:

- Peer-to-peer protocol with no client-server relationship at the CFS layer.
- Three scopes of distribution
  - Logical scope: The distribution occurs within the scope of a VSAN.
  - Physical scope: The distribution spans the entire physical topology.
  - Over a selected set of VSANs: As of Cisco SAN-OS Release 2.1(1a), some applications, such as Inter-VSAN Routing (IVR), require configuration distribution over some specific VSANs. These applications can specify to CFS the set of VSANs over which to restrict the distribution.
- Three modes of distribution
  - Coordinated distributions: Only one distribution is allowed in the fabric at any given time.
  - Uncoordinated distributions: Multiple parallel distributions are allowed in the fabric except when a coordinated distribution is in progress.
  - Unrestricted uncoordinated distributions: As of Cisco SAN-OS Release 2.1(1a), multiple parallel distributions are allowed in the fabric in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.
- Supports a merge protocol that facilitates the merge of application configuration during a fabric merge event (when two independent fabrics merge).

CFS Protocol

The CFS functionality is independent of the lower layer transport. Currently, in Cisco MDS switches, the CFS protocol layer resides on top of the FC2 layer and is peer-to-peer with no client-server relationship. CFS uses the FC2 transport services to send information to other switches. CFS uses a proprietary SW_ILS (0x77434653) protocol for all CFS packets. CFS packets are sent to or from the switch domain controller addresses.

CFS can also use IP to send information to other switches (see the "CFS Distribution over IP" section on page 12-10).

Applications that use CFS are completely unaware of the lower layer transport.

CFS Distribution Scopes

Different applications on the Cisco MDS 9000 Family switches need to distribute the configuration at various levels:

- VSAN level (logical scope)

  Applications that operate within the scope of a VSAN have the configuration distribution restricted to the VSAN. An example application is port security where the configuration database is applicable only within a VSAN.

- Physical topology level (physical scope)

  Applications might need to distribute the configuration to the entire physical topology spanning several VSANs. Such applications include NTP and DPVM (WWN based VSAN), which are independent of VSANs.

- Between two switches

  Applications might only operate between select switches in the fabric. An example application is SCSI Flow Services, which operates between two switches.

CFS Distribution Modes

CFS supports different distribution modes to support different application requirements: coordinated and uncoordinated distributions. Both modes are mutually exclusive. Only one mode is allowed at any given time.

Uncoordinated Distribution

Uncoordinated distributions are used to distribute information that is not expected to conflict with that from a peer. An example is local device registrations such as iSNS. Parallel uncoordinated distributions are allowed for an application.

Coordinated Distribution

Coordinated distributions can have only one application distribution at a given time. CFS uses locks to enforce this. A coordinated distribution is not allowed to start if locks are taken for the application anywhere in the fabric. A coordinated distribution consists of three stages:

1. A fabric lock is acquired.
2. The configuration is distributed and committed.
3. The fabric lock is released.

Coordinated distribution has two variants:

- CFS driven —The stages are executed by CFS in response to an application request without intervention from the application.
- Application driven—The stages are under the complete control of the application.

Coordinated distributions are used to distribute information that can be manipulated and distributed from multiple switches, for example, the port security configuration.

Unrestricted Uncoordinated Distributions

Unrestricted uncoordinated distributions allow multiple parallel distributions in the fabric in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.

Disabling CFS Distribution on a Switch

By default, CFS distribution is enabled. Applications can distribute data and configuration information to all CFS-capable switches in the fabric where the applications exist. This is the normal mode of operation.

You can globally disable CFS on a switch to isolate the applications using CFS from fabric-wide distributions while maintaining physical connectivity. When CFS is globally disabled on a switch, CFS operations are restricted to the switch and all CFS commands continue to function as if the switch were physically isolated.

To globally disable or enable CFS distribution on a switch using Fabric Manager, follow these steps:

Step 1 Choose any CFS feature. For example, expand **Switches > Events** then select **CallHome** in the Physical Attributes pane.

The Information pane shows that feature, with a CFS tab.

Step 2 Click the **CFS** tab to display the CFS state for each switch in the fabric for that feature.

Step 3 Click a value in the Global column—the value changes to a drop-down menu/

Step 4 From the drop-down menu choose **disable** or **enable**

Step 5 Repeat steps 3 and 4 for all switches that you want to disable or enable CFS.

Step 6 Set the Config Action column to **commit.**

Step 7 Click the **Apply Changes** icon to commit the configuration changes for that feature and distribute the changes through CFS, or click the **Undo Changes** icon to discard the changes for that feature.

To globally disable or enable CFS distribution on a switch using Device Manager, follow these steps:

Step 1 Select **Admin > CFS** from the main menu.

You see the CFS dialog box with the CFS status for all features on that switch.

Step 2 Uncheck or check the **Globally Enabled** check box to disable or enable CFS distribution on this switch.

Step 3 Click **Apply** to disable CFS on this switch or **Close** to close the dialog box without making changes.

CFS Application Requirements

All switches in the fabric must be CFS capable. A Cisco MDS 9000 Family switch is CFS capable if it is running Cisco SAN-OS Release 2.0(1b) or later. Switches that are not CFS capable do not receive distributions and result in part of the fabric not receiving the intended distribution.

CFS has the following requirements:

- Implicit CFS usage—The first time you issue a CFS task for a CFS-enabled application, the configuration modification process begins and the application locks the fabric.
- Pending database—The pending database is a temporary buffer to hold uncommitted information. The uncommitted changes are not applied immediately to ensure that the database is synchronized with the database in the other switches in the fabric. When you commit the changes, the pending database overwrites the configuration database (also known as the active database or the effective database).
- CFS distribution enabled or disabled on a per-application basis—The default (enable or disable) for CFS distribution state differs between applications. If CFS distribution is disabled for an application, then that application does not distribute any configuration nor does it accept a distribution from other switches in the fabric.
- Explicit CFS commit—Most applications require an explicit commit operation to copy the changes in the temporary buffer to the application database, to distribute the new database to the fabric, and to release the fabric lock. The changes in the temporary buffer are not applied if you do not perform the commit operation.

Enabling CFS for an Application

All CFS based applications provide an option to enable or disable the distribution capabilities. Features that existed prior to Cisco SAN-OS Release 2.0(1b) have the distribution capability disabled by default and must have distribution capabilities enabled explicitly.

Applications introduced in Cisco SAN-OS Release 2.0(1b) or later have the distribution enabled by default.

The application configuration is not distributed by CFS unless distribution is explicitly enabled for that application.

To enable CFS for a feature using Fabric Manager, follow these steps:

Step 1 Choose a feature on which to enable CFS. For example, expand **Switches > Events** then select **CallHome** in the Physical Attributes pane. The Information pane shows that feature, with a CFS tab. Click the **CFS** tab to display the CFS state for each switch in the fabric for that feature.

Step 2 Decide on which switch(es) to enable CFS. Set the Admin column to either **enable** to enable CFS or **disable** to disable CFS.

Note Enable CFS for all switches in the fabric or VSAN for the feature that uses CFS.

Step 3 Right-click the row you changed to see the pop-up menu. Select **Apply Changes** to apply the CFS configuration change. The CFS tab updates as the CFS changes take effect. Fabric Manager retrieves the status of the CFS change and updates the Last Result column.

To enable CFS for a feature using Device Manager, follow these steps:

Step 1 Click **Admin > CFS** from the main menu.

You see the CFS dialog box with the CFS status for all features on that switch.

Step 2 Decide which feature(s) need CFS. Set the Command column to either **enable** to enable CFS or **disable** to disable CFS.

Note Enable or disable CFS for all switches in the fabric or VSAN for the feature that uses CFS.

Step 3 Click **Pending Differences** to compare the configuration of this feature on this switch to other switches in the fabric or VSAN that have CFS enabled for this feature. Close the Show Pending Diff pop-up.

Step 4 Click **Apply** to apply the CFS configuration change or **Close** to close the dialog box without making changes. Device Manager retrieves the status of the CFS change and updates the Last Command and Result columns.

Locking the Fabric

When you configure (first time configuration) a Cisco SAN-OS feature (or application) that uses the CFS infrastructure, that feature starts a CFS session and locks the fabric. When a fabric is locked, the Cisco SAN-OS software does not allow any configuration changes from a switch, other than the switch holding the lock, to this Cisco SAN-OS feature and issues a message to inform the user about the locked status. The configuration changes are held in a pending database by that application.

If you start a CFS session that requires a fabric lock but forget to end the session, an administrator can clear the session. If you lock a fabric at any time, your user name is remembered across restarts and switchovers. If another user (on the same machine) tries to perform configuration tasks, that user's attempts are rejected.

Committing Changes

A commit operation saves the pending database for all application peers and releases the lock for all switches.

In general, the commit function does not start a session—only a lock function starts a session. However, an empty commit is allowed if configuration changes are not previously made. In this case, a commit operation results in a session that acquires locks and distributes the current database.

When you commit configuration changes to a feature using the CFS infrastructure, you receive a notification about one of the following responses:

- One or more external switches report a successful status—The application applies the changes locally and releases the fabric lock.
- None of the external switches report a successful state—The application considers this state a failure and does not apply the changes to any switch in the fabric. The fabric lock is not released.

You can commit changes for a specified feature by setting the CFS > Config Action to **commit** for that feature.

To commit changes using Fabric Manager for CFS-enabled features, follow these steps:

Step 1 Choose the feature you want to enable CFS for. For example, expand **Switches** expand **Events** then select **CallHome** from the Physical Attributes pane. The Information pane shows that feature, with a CFS tab. Click the **CFS** tab to display the CFS state for each switch in the fabric for that feature.

Step 2 Right-click the value in the Config Action column for any switch and select an option from the drop-down menu (Copy, Paste, Export to File, Print Table, Detach Table).

Step 3 Click the **Apply Changes** icon to commit the configuration changes for that feature and distribute the changes through CFS, or click the **Undo Changes** icon to discard the changes

Fabric Manager retrieves the status of the CFS change and updates the Last Command and Last Result columns for the feature or VSAN.

To commit changes using Device Manager for CFS-enabled features, follow these steps:

Step 1 Click **Admin** and select **CFS** from the main menu. You see the CFS dialog box with the CFS status for all features on that switch.

Step 2 For each applicable feature, set the Command column to **commit** to commit the configuration changes for that feature and distribute the changes through CFS, or set it to **abort** to discard the changes for that feature and release the fabric lock for CFS for that feature.

Step 3 Optionally, provide a **Type** or **VsanID** as the basis for the CFS distribution for CFS features that require this.

Step 4 Click **Pending Differences** to check the configuration of this feature on this switch as compared to other switches in the fabric or VSAN that have CFS enabled for this feature.

Step 5 Click **Apply** to apply the CFS configuration change or **Close** to close the dialog box without making changes.

Device Manager retrieves the status of the CFS change and updates the Last Command and Result columns.

Caution If you do not commit the changes, they are not saved to the running configuration.

Discarding Changes

If you discard configuration changes, the application flushes the pending database and releases locks in the fabric. Both the abort and commit functions are only supported from the switch from which the fabric lock is acquired.

You can discard changes for a specified feature by setting the Command column value to **disable** for that feature then clicking **Apply**.

Saving the Configuration

Configuration changes that have not been applied yet (still in the pending database) are not shown in the running configuration. The configuration changes in the pending database overwrite the configuration in the effective database when you commit the changes.

Caution If you do not commit the changes, they are not saved to the running configuration.

The CISCO-CFS-MIB contains SNMP configuration information for any CFS-related functions. Refer to the *Cisco MDS 9000 Family MIB Quick Reference* for more information on this MIB.

Clearing a Locked Session

You can clear locks held by an application from any switch in the fabric. This option is provided to rescue you from situations where locks are acquired and not released. This function requires Admin permissions.

Caution Exercise caution when using this function to clear locks in the fabric. Any pending configurations in any switch in the fabric is flushed and lost.

CFS Merge Support

An application keeps the configuration synchronized in a fabric through CFS. Two such fabrics might merge as a result of an ISL coming up between them. These two fabrics could have two different sets of configuration information that need to be reconciled in the event of a merge. CFS provides notification each time an application peer comes online. If a fabric with M application peers merges with another fabric with N application peers and if an application triggers a merge action on every such notification, a link-up event results in M*N merges in the fabric.

CFS supports a protocol that reduces the number of merges required to one by handling the complexity of the merge at the CFS layer. This protocol runs per application per scope. The protocol involves selecting one switch in a fabric as the merge manager for that fabric. The other switches do not play any role in the merge process.

During a merge, the merge manager in the two fabrics exchange their configuration databases with each other. The application on one of them merges the information, decides if the merge is successful, and informs all switches in the combined fabric of the status of the merge.

In case of a successful merge, the merged database is distributed to all switches in the combined fabric and the entire new fabric remains in a consistent state. You can recover from a merge failure by starting a distribution from any of the switches in the new fabric. This distribution restores all peers in the fabric to the same configuration database.

Displaying CFS Configuration Information

To display the status of CFS distribution on the switch using Device Manager, follow these steps:

Step 1 Click **Admin > Cisco Fabric Services (CFS)**.

You see the CFS dialog box. This dialog box displays the distribution status of each feature using CFS, which currently registered applications are using CFS, and the result of the last successful merge attempt.

Step 2 Select a row and click **Details** to view more information about the feature.

CFS Distribution over IP

You can configure CFS to distribute information over IP for networks containing switches that are not reachable over Fibre Channel. CFS distribution over IP supports the following features:

- Physical distribution over an entirely IP network.
- Physical distribution over a hybrid Fibre Channel and IP network with the distribution reaching all switches that are reachable over either Fibre Channel or IP

Note The switch attempts to distribute information over Fibre Channel first and then over the IP network if the first attempt over Fibre Channel fails. CFS does not send duplicate messages if distribution over both IP and Fibre Channel is enabled.

- Distribution over IP version 4 (IPv4) or IP version 6 (IPv6).

Note CFS cannot distribute over both IPv4 and IPv6 from the same switch.

- Keep-alive mechanism to detect network topology changes using a configurable multicast address.
- Compatibility with Cisco MDS SAN-OS Release 2.x.
- Distribution for logical scope applications is not supported because the VSAN implementation is limited to Fibre Channel.

Figure 12-1 shows a network with both Fibre Channel and IP connections. Node A forwards an event to node B over Fibre Channel. Node B forwards the event node C and node D using unicast IP. Node C forwards the event to node E using Fibre Channel.

Figure 12-1 Network Example 1 with Fibre Channel and IP Connections

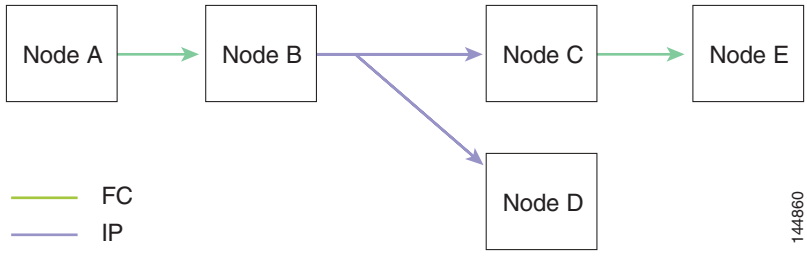

Figure 12-2 is the same as Figure 12-1 except that node C and node D are connected using Fibre Channel. All processes is the same in this example because node B has node C and node D the distribution list for IP. Node C does not forward to node D because node D is already in the distribution list from node B.

Figure 12-2 Network Example 2 with Fibre Channel and IP Connections

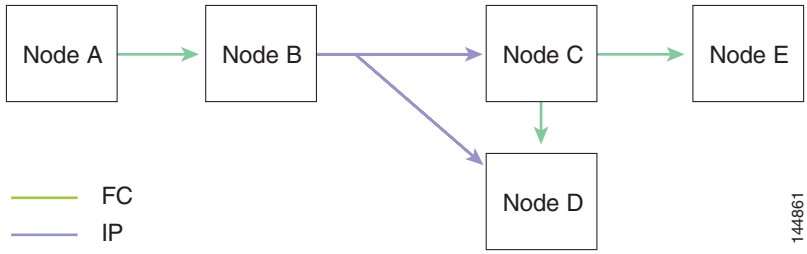

Figure 12-3 is the same as Figure 12-2 except that node D and node E are connected using IP. Both node C and node D forward the event to E because the node E is not in the distribution list from node B.

Figure 12-3 Network Example 3 with Fibre Channel and IP Connections

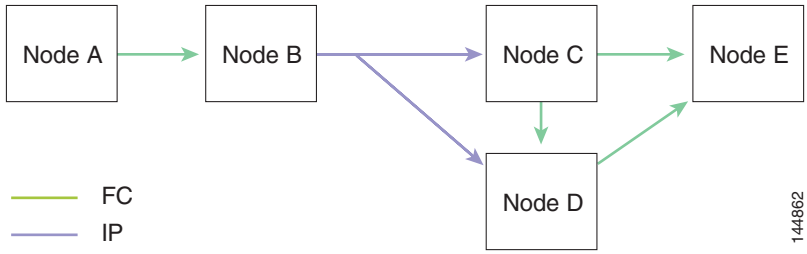

Note CFS has been assigned port 7456 by the Internet Assigned Numbers Authority (IANA). CFS uses this port for both multicast and unicast message distribution.

Note CFS uses only the management interface for distribution over an IP network.

A CFS Example Using Fabric Manager

This procedure is an example of what you see when you use Fabric Manager to configure a feature that uses CFS. For specific procedures for features that use CFS, refer to that feature's documentation.

To configure a feature that uses CFS using Fabric Manager, follow these steps:

Step 1 Select the CFS-capable feature you want to configure. For example, expand **Switches** expand **Clock** then select **NTP** in the Physical Attributes pane. You see the feature configuration in the Information pane, including a CFS tab.

Step 2 Click the **CFS** tab.

You see the CFS configuration and status for each switch.

Step 3 Right-click a value in the Admin column then choose **enable** for a switch.

Step 4 Repeat step 3 for all switches in the fabric.

Note A warning displays if you do not enable CFS for all switches in the fabric for this feature.

Step 5 Check the **Master** check box for the switch to act as the merge master for this feature.

Step 6 Click the value in the Config Action column then choose **commit** from the drop-down list for each switch that you enabled for CFS.

Step 7 Click the **Servers tab** in the Information pane. You see the configuration for this feature based on the master switch.

Step 8 Modify the feature configuration. For example, right-click the name in the Master column and select the **Create Row** icon to create a server for NTP.

a. Set the ID, and the Name or IP Address for the NTP server.

b. Set the **Mode** radio button and optionally check the Preferred check box.

c. Click **Create** to add the server.
Fabric Manager sends the request to the master switch and updates the Last Result column under the CFS tab with the "pending" status.

Step 9 Click the **CFS** tab.

Step 10 Set the Config Action column to **Apply Changes** to distribute the feature change through the fabric. Fabric Manager only changes the status of the Config View as column to "running" when **commit**, **clear,** or **abort** is selected and applied.

Note Fabric Manager does not change the status to "pending" if **enable** is selected, because the "pending" status does not apply until the first actual change is made.

Step 11 Click the **Apply Changes** icon to commit the configuration changes for that feature and distribute the changes through CFS, or click the **Undo Changes icon** to discard the changes for that feature.

Note When using CFS with features like DPVM and device alias, you must select **commit** at the end of each configuration. If the session is locked, you must exit the feature by selecting **abort**.

To configure the master or seed switch for distribution for each feature using Fabric Manager, follow these steps:

Step 1 Choose the feature that needs a merge master for CFS. For example, expand **Switches** expand **Events** and select **CallHome** from the Physical Attributes pane. The Information pane shows that feature including a CFS tab. Click the **CFS** tab to display the CFS state for each switch in the fabric for that feature.

Step 2 Check the Master column check box for the switch to act as the merge master for this feature.

Step 3 Click the **Apply Changes** icon to select this switch as master for future CFS distributions, or click the **Undo changes** icon to discard your unsaved changes.

A CFS Example Using Device Manager

This procedure is an example of what you see when you use Device Manager to configure a feature that uses CFS. For specific procedures for features that use CFS, refer to that feature's documentation.

To configure a feature that uses CFS using Device Manager, follow these steps:

Step 1 Open the dialog box for any CFS-capable feature. Device Manager checks to see whether CFS is enabled. It also checks to see if there is a lock on the feature by checking for at least one entry in the Owner table. If CFS is enabled and there is a lock, Device Manager sets the status to "pending" for that feature. You see a dialog box displaying the lock information.

Step 2 Click **Continue** or **Cancel** when prompted. If you continue, Device Manager remembers the CFS status.

Step 3 Select **Admin > CFS (Cisco Fabric Services)** to view the user name of the CFS lock holder. Click the locked feature and click **Details**.

Step 4 Click the **Owners** tab and look in the UserName column.

Note Device Manager does not monitor the status of the feature across the fabric until you click **Refresh**. If a user on another CFS-enabled switch attempts to configure the same feature, they do not see the "pending" status. However, their configuration changes are rejected by your switch.

Step 5 If CFS is enabled and there is no lock, Device Manager sets the status to running for that feature. You then see a dialog box for the feature. As soon as you perform a creation, deletion, or modification, Device Manager changes the status to "pending" and displays the updated information from the pending database.

Step 6 View the CFS table for a feature. Device Manager only changes the status to "running" when **commit**, **clear,** or **abort** is selected and applied. Device Manager will not change the status to "pending" if **enable** is selected, because the "pending" status does not apply until the first actual change is made.

The Last Command and Result fields are blank if the last command is **noOp**.

Note When using CFS with features like DPVM and device alias, you must select **commit** at the end of each configuration. If the session is locked, you must exit the feature by selecting **abort**.

Default Settings

Table 12-1 lists the default settings for CFS configurations.

Table 12-1 Default CFS Parameters

Parameters	Default
CFS distribution	Enabled.
Database changes	Implicitly enabled with the first configuration change.
Application distribution	Differs based on application.
Commit	Explicit configuration is required.
CFS over IP	Disabled.
IPv4 multicast address	226.68.71.84
IPv6 multicast address	ff02::e244:4754

CHAPTER 13

Software Images

This chapter describes how to install and upgrade Cisco MDS SAN-OS software images. It includes the following sections:

About Software Images

Each switch is shipped with a Cisco MDS SAN-OS operating system for Cisco MDS 9000 Family switches. The Cisco MDS SAN-OS consists of two images—the kickstart image and the system image. To upgrade the switch to a new image, you must specify the variables that direct the switch to the images.

- To select the kickstart image, use the KICKSTART variable.
- To select the system image, use the SYSTEM variable.

The images and variables are important factors in any install procedure. You must specify the variable and the image to upgrade your switch. Both images are not always required for each install.

Note Unless explicitly stated, the software install procedures in this section apply to any switch in the Cisco MDS 9000 Family.

Dependent Factors for Software Installation

The software image install procedure is dependent on the following factors:

- Software images—The kickstart and system image files reside in directories or folders that can be accessed from the Cisco MDS 9000 Family switch prompt.
- Image version—Each image file has a version.
- Flash disks on the switch—The bootflash: resides on the supervisor and the CompactFlash disk is inserted into the slot0: device.
- Supervisor modules—There are single or dual supervisor modules.

Selecting the Correct Software Images for Cisco MDS 9500 Series Switches

The Supervisor-1 and Supervisor-2 modules supported by Cisco MDS 9500 Series switches require different system and kickstart images. You can determine which images to use on your switch by the naming conventions shown in Table 13-1.

Table 13-1 Supervisor Module Software Image Naming Conventions

Cisco MDS 9500 Series Switch Type	Supervisor Module Type	Naming Convention
9506 or 9509	Supervisor-1 module	Filename begins with m9500-sf1ek9
	Supervisor-2module	Filename begins with m9500-sf2ek9
9513	Supervisor-2 module	Filename begins with m9500-sf2ek9

Essential Upgrade Prerequisites

Before attempting to migrate to any software image version, follow these guidelines:

- Customer Service

  Before performing any software upgrade, contact your respective customer service representative to review your software upgrade requirements and to provide recommendations based on your current operating environment.

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

- Scheduling

  Schedule the upgrade when the fabric is stable and steady. Ensure that everyone who has access to the switch or the network is not configuring the switch or the network during this time. All configurations are disallowed at this time.

- Space

  Verify that sufficient space is available in the location where you are copying the images. This location includes the active and standby supervisor bootflash: (internal to the switch).

  - Standby supervisor module bootflash: file system (see the Chapter 11, "Initial Configuration").
  - Internal bootflash offers approximately 200 MB of user space.

- Hardware

  Avoid power interruption during any install procedure. These kinds of problems can corrupt the software image.

- Connectivity (to retrieve images from remote servers)

  - Configure the IPv4 address or IPv6 address for the 10/100/1000 BASE-T Ethernet port connection (interface mgmt0).

Note 1000 BASE-T Ethernet is only available on Supervisor-2 modules.

  - Ensure the switch has a route to the remote server. The switch and the remote server must be in the same subnetwork if you do not have a router to route traffic between subnets.

- Images

  - Ensure that the specified system and kickstart images are compatible with each other.
  - If the kickstart image is not specified, the switch uses the current running kickstart image.
  - If you specify a different system image, ensure that it is compatible with the running kickstart image.
  - Retrieve images in one of two ways:

    Locally—images are locally available on the switch.

    Remotely—images are in a remote location and the user specifies the destination using the remote server parameters and the file name to be used locally.

- Terminology

  Table 13-2 summarizes terms used in this chapter with specific reference to the install and upgrade process.

Table 13-2 Terms Specific to This Chapter

Term		Definition
bootable		The modules ability to boot or not boot based on image compatibility.
impact		The type of software upgrade mechanism—disruptive or nondisruptive.
install-type	reset	Resets the module.
	sw-reset	Resets the module immediately after switchover.
	rolling	Upgrades each module in sequence.
	copy-only	Updates the software for BIOS, loader, or bootrom.

- Tools
  - Verify connectivity to the remote server by clicking **Verify Remote Server** in the Software Install Wizard in Fabric Manager.
  - Ensure that the required space is available for the image files to be copied by selecting **Physical > System** on Device Manager and checking the NVRAM statistics.
  - We recommend the Software Install Wizard in Fabric Manager to upgrade your software. This wizard upgrades all modules in any Cisco MDS 9000 Family switch (see the "Benefits of Using the Software Install Wizard" section on page 13-6).
  - Run only one installation on a switch at any time.
  - Do not issue another command while running the installation.
  - Do the installation on the active supervisor module, not the standby supervisor module.

Note If the switching module(s) are not compatible with the new supervisor module image, some traffic disruption may be noticed in the related modules, depending on your configuration. These modules are identified in the summary when you use the Installation Wizard. You can choose to proceed with the upgrade or end at this point.

Note The Software Install Wizard displays a summary of changes that are made to your configuration and waits for your authorization to continue the installation.

Note To preserve the FC IDs in your configuration, verify that the persistent FC ID feature is enabled before rebooting. This feature is enabled by default. In earlier releases, the default is disabled. See the "FC IDs" section on page 22-15

Software Upgrade Methods

You can upgrade software without any disruptions using the Cisco MDS SAN-OS software designed for mission-critical high availability environments. To realize the benefits of nondisruptive upgrades on the Cisco MDS 9500 Directors, we highly recommend that you install dual supervisor modules.

You can upgrade any switch in the Cisco MDS 9000 Family using one of the following methods:

You can upgrade any switch in the Cisco MDS 9000 Family using one of the following methods:

- Automatic - you can use the Fabric Manager Software Install Wizard for Cisco MDS SAN-OS switches as described in the "Using the Software Install Wizard" section on page 13-7.
- Manual - for information on manual upgrades, see the *Cisco MDS 9000 Family CLI Configuration Guide* or the *Cisco MDS 9020 Switch Configuration Guide and Command Reference*.

In some cases, regardless of which process you use, the software upgrades may be disruptive. These exception scenarios can occur under the following conditions:

- A single supervisor system with kickstart or system image changes.
- A dual supervisor system with incompatible system software images.

Determining Software Compatibility

If the running image and the image you want to install are incompatible, the software reports the incompatibility. In some cases, you may decide to proceed with this installation. If the active and the standby supervisor modules run different versions of the image, both images may be HA compatible in some cases and incompatible in others.

Compatibility is established based on the image and configuration:

- Image incompatibility—The running image and the image to be installed are not compatible.
- Configuration incompatibility—There is a possible incompatibility if certain features in the running image are turned off as they are not supported in the image to be installed. The image to be installed is considered incompatible with the running image if one of the following statements is true:
  - An incompatible feature is enabled in the image to be installed and it is not available in the running image and may cause the switch to move into an inconsistent state. In this case, the incompatibility is *strict*.
  - An incompatible feature is enabled in the image to be installed and it is not available in the running image and does not cause the switch to move into an inconsistent state. In this case, the incompatibility is *loose*.

Tip The Software Install Wizard compares and presents the results of the compatibility before proceeding with the installation. You can exit if you do not want to proceed with these changes.

Automated Upgrades

The Software Install Wizard upgrades all modules in any Cisco MDS 9000 Family switch. Figure 13-1 provides an overview of the switch status before and after using Software Install Wizard.

Figure 13-1 The Effect of the Software Install Wizard

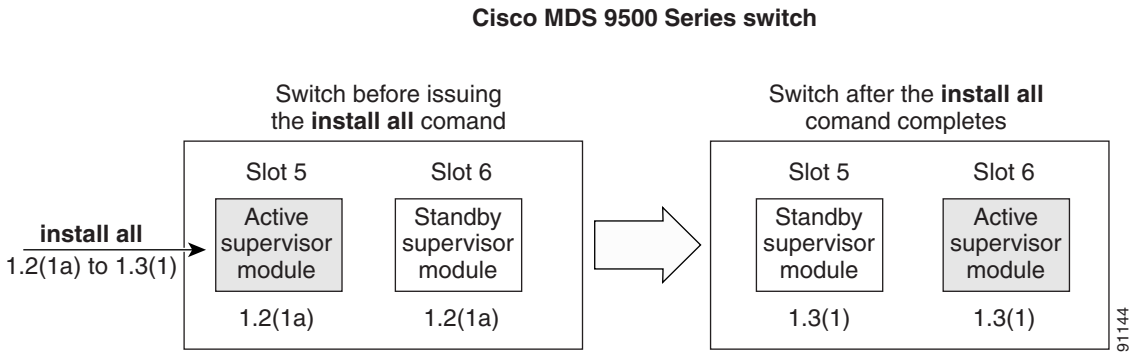

The Software Install Wizard automatically verifies if the standby supervisor module is functioning (if present). If it is not functioning, it reloads that module and uses the **force download** option to force it to function.

Benefits of Using the Software Install Wizard

The Software Install Wizard provides the following benefits:

- You can upgrade the entire switch using just one procedure.
- You can receive descriptive information on the intended changes to your system before you continue with the installation.
- You have the option to cancel the installation.
- You can upgrade the entire switch using the least disruptive procedure.
- You can see the progress of this command on the console, Telnet, and SSH screens:
  - After a switchover process, you can see the progress from both the supervisor modules.
  - Before a switchover process, you can only see the progress from the active supervisor module.
- The Software Install Wizard automatically checks the image integrity. This includes the running kickstart and system images.
- The Software Install Wizard performs a platform validity check to verify that a wrong image is not used—for example, to check if an MDS 9500 Series image is used inadvertently to upgrade an MDS 9200 Series switch.
- After issuing the installation, if any step in the sequence fails, the wizard completes the step in progress and ends.

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

Recognizing Failure Cases

The following situations cause the installation to end:

- If the standby supervisor module bootflash: file system does not have sufficient space to accept the updated image.
- If the specified system and kickstart images are not compatible.
- If the fabric or switch is configured while the upgrade is in progress.
- If a module is removed while the upgrade is in progress.
- If the switch has any power disruption while the upgrade is in progress.
- If the entire path for the remote location is not specified accurately.
- If images are incompatible after an upgrade. For example, a switching module image may be incompatible with the system image, or a kickstart image may be incompatible with a system image. This is also identified by the Software Install Wizard compatibility check.

Caution If the installation is ended, be sure to verify the state of the switch at every stage and reissue the command after 10 seconds. If you reissue the installation within the 10-second span, it is rejected with an error message indicating that an installation is currently in progress.

Tip Most configurations are disallowed while the installation is in progress. However, configurations coming through the CFS applications are allowed and may affect the upgrade procedure.

Using the Software Install Wizard

You can use the Software Install Wizard to install Cisco SAN-OS images on supported switches.

Note The Software Install Wizard supports installation and upgrade for Cisco MDS 9020 Fabric Switch or Cisco FabricWare. For successful installation and upgrade specify the TFTP server address that the Cisco MDS 9020 Fabric Switch should use.

Note Before you use this wizard, be sure the standby supervisor management port is connected.

To use the Software Install Wizard, follow these steps:

Step 1 Click the Software Install Wizard icon in the toolbar (see Figure 13-2).

Figure 13-2 Software Install Wizard Icon

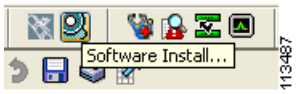

You see the Software Install Wizard.

Step 2 Select the switches you want to install images on. You must select at least one switch to proceed. When finished, click **Next**.

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

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

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

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

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

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

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

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

Step 9 Optionally, check **Ignore version check results** to bypass a version check.

Note The version check provides information about the impact of the upgrade for each module on the switch. It also shows any HA-related incompatibilities that might result. You see a final dialog box at this stage, prompting you to confirm that this check should be performed. We recommend that you do not ignore the version check results.

Caution If **Ignore version check results** is checked, the upgrade will proceed even if the current switch version is newer than the version you are installing.

Step 10 Click **Finish** to start the installation or click **Cancel** to leave the installation wizard without installing new images.

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

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

Upgrading Services Modules

Any Fibre Channel switching module supports nondisruptive upgrades. The14/2-port Multiprotocol Services (MPS-14/2)) module supports nondisruptive upgrades for the Fibre Channel ports. Any software upgrade for the two Gigabit Ethernet ports in this module is disruptive. See Chapter 47, "Configuring IP Storage" for further information on MPS-14/2 modules.

Caution Any software upgrade for the Caching Services Module (CSM) and the IP Storage (IPS) services modules is disruptive.

CSMs and IPS modules use a rolling upgrade install mechanism to guarantee a stable state for each module in the switch:

- Each IPS module in a switch requires a 5-minute delay before the next IPS module is upgraded. See Chapter 47, "Configuring IP Storage" for more information on IPS modules.
- Each CSM module requires a 30-minute delay before the next CSM module is upgraded. See the *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for more information on CSMs.

Maintaining Supervisor Modules

This section includes general information about replacing and using supervisor modules effectively.

This section includes the following topics:

Replacing Supervisor Modules

To avoid packet loss when removing a supervisor module from a Cisco MDS 9500 Series Director, take the supervisor modules out of service before removing the supervisor module.

Note You must remove and reinsert or replace the supervisor module to bring it into service.

Migrating from Supervisor-1 Modules to Supervisor-2 Modules

As of Cisco MDS SAN-OS Release 3.0(1) and later, the Cisco MDS 9509 and 9506 Directors support both Supervisor-1 and Supervisor-2 modules. Supervisor-1 and Supervisor-2 modules cannot be installed in the same switch, except during migration. Both the active and standby supervisor modules must be of the same type, either Supervisor-1 or Supervisor-2 modules. For Cisco MDS 9513 Directors, both supervisor modules must be Supervisor-2 modules.

The procedure described in this section ensures that your configuration is correctly synchronized after completing the migration.

Caution Migrating your supervisor modules is a disruptive operation.

Note Migrating from Supervisor-2 modules to Supervisor-1 modules is not supported.

To migrate from using Supervisor-1 modules to Supervisor-2 modules on a Cisco MDS 9509 or 9506 Director, follow these steps:

Step 1 Ensure that the configured domain ID is the same as the current domain ID for every VSAN on the switch by following these steps:

a. Issue a **show vsan** command to display all the VSANs on the switch.

```
switch# show vsan
vsan 1 information
         name:VSAN0001  state:active
         interoperability mode:default
         loadbalancing:src-id/dst-id/oxid
         operational state:down

vsan 2 information
         name:VSAN0002  state:active
         interoperability mode:default
         loadbalancing:src-id/dst-id/oxid
         operational state:down

vsan 10 information
         name:VSAN0010  state:active
         interoperability mode:default
         loadbalancing:src-id/dst-id
         operational state:down

vsan 4094:isolated_vsan
```

b. Display the current and configured domain IDs for a VSAN.

```
switch# show fcdomain vsan 1
The local switch is the Principal Switch.

Local switch run time information:
        State: Stable
        Local switch WWN:    20:01:00:05:30:00:35:df
        Running fabric name: 20:01:00:05:30:00:35:df
        Running priority: 128
        Current domain ID: 0x6a(106)

Local switch configuration information:
        State: Enabled
        FCID persistence: Enabled
        Auto-reconfiguration: Disabled
        Contiguous-allocation: Disabled
        Configured fabric name: 20:01:00:05:30:00:28:df
        Configured priority: 128
        Configured domain ID: 0x00(0) (preferred)

Principal switch run time information:
        Running priority: 128
```

c. Change the configured domain ID if it differs from the current domain ID.

```
switch# config t
switch(config)# fcdomain domain 106 static vsan 1
switch(config)# exit
switch#
```

d. Repeat Step b and Step c for each VSAN on the switch.

Step 2 Save the configuration.

```
switch# copy running-config startup-config
```

Step 3 Verify that the switch is running Cisco SAN-OS Release 3.0(1) or later. Upgrade the switch, if necessary (see the "Automated Upgrades" section on page 13-5).

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

Software
  BIOS:      version 0.0.11
  kickstart: version 3.0(1) [build 3.0(0.294)] [gdb]
  system:    version 3.0(1) [build 3.0(0.294)] [gdb]
...
```

Step 4 Issue a **show module** command to determine which Supervisor-1 module is the standby.

```
switch# show module
Mod  Ports  Module-Type                      Model              Status
---  -----  -------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module               DS-X9016           ok
2    32     Storage Services Module          DS-X9032-SSM       ok
3    8      IP Storage Services Module       DS-X9308-SMIP      ok
4    12     1/2/4 Gbps FC Module             DS-X9112           ok
5    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    ha-standby
6    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    active *
...
```

Step 5 Take the standby Supervisor-1 module out of service.

```
switch# out-of-service module 6
```

Step 6 Verify that the standby Supervisor-1 module is powered down.

```
switch# show module
Mod  Ports  Module-Type                      Model              Status
---  -----  -------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module               DS-X9016           ok
2    32     Storage Services Module          DS-X9032-SSM       ok
3    8      IP Storage Services Module       DS-X9308-SMIP      ok
4    12     1/2/4 Gbps FC Module             DS-X9112           ok
5    0      Supervisor/Fabric-1                                 powered-dn
6    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    active *
...
```

Step 7 Remove the standby Supervisor-1 module from the chassis.

Step 8 Install the Supervisor-2 module in the chassis.

Step 9 Establish a console session on the standby Supervisor-2 module console port.

Step 10 If the `loader>` prompt appears on the standby Supervisor-2 module console session, perform the following steps. Otherwise continue to Step 11.

a. Verify the Cisco SAN-OS system image and kickstart image on the standby Supervisor-2 module bootflash:.

```
loader> dir bootflash:
40295206     Aug 05 15:23:51 1980  ilc1.bin
12456448     Jul 30 23:05:28 1980  kickstart-image
12288        Jun 23 14:58:44 1980  lost+found/
27602159     Jul 30 23:05:16 1980  system-image
12447232     Aug 05 15:08:30 1980  kickstart-image2
28364853     Aug 05 15:11:57 1980  system-image2

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

b. If the images are present boot the standby Supervisor-2 module skip to Step h. Otherwise, continue to the next step.

c. Enter the local IP address and the subnet mask for the switch, and press **Enter**.

```
loader> ip address 10.16.1.2 255.255.255.0
Found Intel EtherExpressPro100 82559ER at 0xe800, ROM address 0xc000
Probing...[Intel EtherExpressPro100 82559ER]Ethernet addr: 00:05:30:00:52:27
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 0.0.0.0
Gateway: 0.0.0.0
```

d. Enter the IP address of the default gateway, and press **Enter**.

```
loader> ip default-gateway 10.16.1.1
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 0.0.0.0
Gateway: 172.16.1.1
```

e. Boot the kickstart image file from the bootflash: (if present) or from a server.

```
loader> boot tftp://10.16.10.100/kickstart-latest
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 172.16.10.100
Gateway: 172.16.1.1
Booting: /kick-282 console=ttyS0,9600n8nn quiet loader_ver= "1.0(2)"....
.........................................Image verification OK
Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..... done.
Loading system software
INIT: Sending processes the TERM signal
Sending all processes the TERM signal... done.
Sending all processes the KILL signal... done.
Entering single-user mode...
INIT: Going single user
INIT: Sending processes the TERM signal
switch(boot)#
```

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

f. Download a Cisco SAN-OS system image to the Supervisor-2 module from a TFTP server.

```
switch(boot)# copy tftp://10.16.10.100/system-img bootflash:system-img
Trying to connect to tftp server......
```

g. Download a kickstart image to the Supervisor-2 module from a TFTP server, if necessary.

```
switch(boot)# copy tftp://10.16.10.100/kickstart-img bootflash:kickstart-img
Trying to connect to tftp server......
```

h. Boot the standby Supervisor-2 module.

```
loader> boot bootflash:kickstart-imag bootflash:system-img
```

Step 11 Verify that the standby Supervisor-2 module is in warm standby state using a **show system redundancy status** command on the active Supervisor-1 module session.

```
switch# show system redundancy status
Redundancy mode
---------------
      administrative:   HA
         operational:   Warm

This supervisor (sup-2)
-----------------------
    Redundancy state:   Active
    Supervisor state:   Active
      Internal state:   Active with warm standby

Other supervisor (sup-1)
------------------------
    Redundancy state:   Standby
    Supervisor state:   Warm standby
      Internal state:   Warm standby
```

Step 12 Copy the running configuration to the startup configuration on the active Supervisor-1 module to ensure that any running configuration changes are saved to the startup configuration and the ASCII configuration is synchronized and up to date on the warm standby Supervisor-2 module.

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

Step 13 If your switch has SSMs installed and intelligent services are provisioned, perform the following steps. Otherwise, continue to Step 14.

a. Power down all SSMs on the switch.

```
switch# config t
switch(config)# poweroff module 2
switch(config)# exit
switch#
```

Caution Do not copy the running configuration to the startup configuration after powering down the SSMs. If you do, you will lose the configuration on the SSM interfaces.

b. Verify that the SSMs are powered down.

```
switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module                DS-X9016           ok
2    32     Storage Services Module                              powered-dn
3    8      IP Storage Services Module        DS-X9308-SMIP      ok
4    12     1/2/4 Gbps FC Module              DS-X9112           ok
5    0      Supervisor/Fabric-2               DS-X9530-SF2-K9    ha-standby
6    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
...
```

c. Copy the contents of the SSM NVRAM to the standby Supervisor 2 module.

```
switch# copy ssm-nvram standby-sup
```

Step 14 Initiate a switchover on the active Supervisor-1 module to power it down and cause the standby Supervisor-2 module to become the active supervisor module.

```
switch# system switchover
```

Step 15 Verify that the Supervisor-1 module is powered down.

```
switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module                DS-X9016           ok
2    32     Storage Services Module           DS-X9032-SSM       ok
3    8      IP Storage Services Module        DS-X9308-SMIP      ok
4    12     1/2/4 Gbps FC Module              DS-X9112           ok
5    0      Supervisor/Fabric-2               DS-X9530-SF2-K9    active *
6    0      Supervisor/Fabric-1                                  powered-dn
...
```

Step 16 Remove the Supervisor-1 module from the chassis.

Step 17 Set the baud rate on the active Supervisor-2 module console session to the default value of 9600.

```
switch# config t
switch(config)# line console
switch(config-console)# speed 9600
switch(config-console)# end
switch# show line console
line Console:
    Speed:        9600 bauds
    Databits:     8 bits per byte
    Stopbits:     1 bit(s)
    Parity:       none
    Modem In: Disable
    Modem Init-String -
        default : ATQ0V1H0S0=1\015
```

Step 18 Install the other Supervisor-2 module in the chassis.

Step 19 Verify that the standby Supervisor-2 module is in the HA-standby state.

```
switch# show system redundancy status
Redundancy mode
---------------
      administrative:   HA
         operational:   HA

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

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

Step 20 If the Cisco MDS SAN-OS system image on the supervisor modules is the desired release, issue the **install all** command.

```
switch# install all
```

If you want a different release of the Cisco SAN-OS system image running on the switch, issue the **install all** command specifying the system image to perform a hitless upgrade (see the "Automated Upgrades" section on page 13-5).

```
switch# install all system tftp://10.16.10.100/system-img
```

Standby Supervisor Boot Variable Version

If the standby supervisor module boot variable images are not the *same* version as those running on the active supervisor module, the software forces the standby supervisor module to run the same version as the active supervisor module.

If you specifically set the boot variables of the standby supervisor module to a different version and reboot the standby supervisor module, the standby supervisor module will only load the specified boot variable if the same version is also running on the active supervisor module. At this point, the standby supervisor module is *not* running the images set in the boot variables.

Standby Supervisor Bootflash Memory

When updating software images on the standby supervisor module, verify that there is enough space available for the image It is a good practice to remove older versions of Cisco MDS SAN-OS images and kickstart images.

To verify the space on the standby supervisor using Device Manager, follow these steps:

Step 1 Click **Admin** > **Flash Files**.

Step 2 Select the standby supervisor from the Partition drop-down list.

At the bottom of the Flash Files dialog box, you see the space used and free space.

Standby Supervisor Boot Alert

If a standby supervisor module fails to boot, the active supervisor module detects that condition and generates a Call Home event and a system message and reboots the standby supervisor module approximately 3 to 6 minutes after the standby supervisor module moves to the `loader>` prompt.

The following system message is issued:

```
%DAEMON-2-SYSTEM_MSG:Standby supervisor failed to boot up.
```

This error message is also generated if one of the following situations apply:

- You remain at the `loader>` prompt for an extended period of time.
- You do not set the boot variables appropriately.

Replacing Modules

When you replace any module (supervisor, switching, or services module), you must ensure that the new module is running the same software version as the rest of the switch.

Refer to *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for configuration details on replacing the caching services module (CSM).

Note When a spare standby supervisor module is inserted, it uses the same image as the active supervisor module. The Cisco SAN-OS software image is not automatically copied to the standby flash device.

Tip Use the Software Install Wizard to copy the Cisco SAN-OS software image to the standby supervisor bootflash device.

Using the Software Install Wizard after replacing any module, ensures the following actions:

- The proper system and kickstart images are copied on the standby bootflash: file system.
- The proper boot variables are set.
- The loader and the BIOS are upgraded to the same version available on the active supervisor module.

To replace a module in any switch in the Cisco MDS 9200 Series or 9500 Series using Device Manager, follow these steps:

Step 1 Create a backup of your existing configuration file, if required, by clicking **Admin > Copy Configuration** and selecting **runningConfig** to **startupConfig**.

Step 2 Replace the required module as specified in the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*.

Step 3 Verify that space is available on the standby supervisor bootflash by clicking **Admin > Flash Files** and selecting the **sup-standby**. It is a good practice to remove older versions of Cisco MDS SAN-OS images and kickstart images.

Step 4 Use the Software Install Wizard to ensure that the new module is running the same software as the rest of the switch.

Step 5 Wait until the new module is online and then ensure that the replacement was successful by clicking **Physical > Modules** in Device Manager.

Default Settings

Table 13-3 lists the default image settings for all Cisco MDS 9000 Family switches.

Table 13-3 Default Image Settings

Parameters	Default
Kickstart image	No image is specified.
System image	No image is specified.

CHAPTER 14

Managing Configuration Files

This chapter describes how to initially configure switches using the configuration files so they can be accessed by other devices. This chapter includes the following sections:

About Flash Devices

Every switch in the Cisco MDS 9000 Family contains one internal bootflash (see Figure 14-1). The Cisco MDS 9500 Series additionally contains one external CompactFlash called slot0 (see Figure 14-1 and Figure 14-2).

Figure 14-1 Flash Devices in the Cisco MDS 9000 Supervisor Module

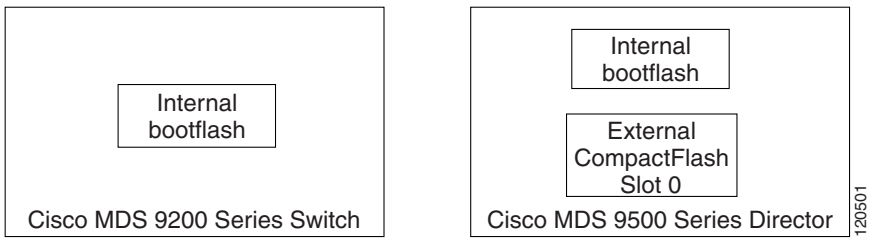

Figure 14-2 External CompactFlash in the Cisco MDS 9000 Supervisor Module

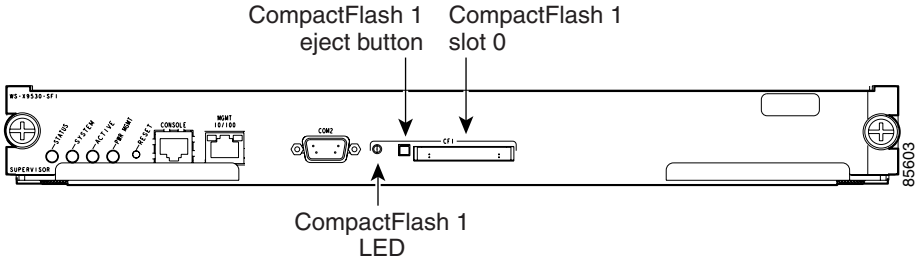

Internal bootflash:

All switches in the Cisco MDS 9000 Family have one internal bootflash: that resides in the supervisor or switching module.You have access to two locations within the internal bootflash: file system.

- The volatile: file system provides temporary storage, and it is also the default location for file system commands. Files in temporary storage (volatile:) are erased when the switch reboots.
- The bootflash: (nonvolatile storage) file system provides permanent storage. The files in bootflash: are preserved through reboots and power outages.

Formatting Flash Devices and File Systems

By formatting a Flash devices or a file system, you are clearing out the contents of the device or the file system and restoring it to its factory-shipped state.

See the "About Flash Devices" section on page 14-1 and the "Using the File System" section on page 14-2.

Initializing Internal bootflash:

When a switch is shipped, the switch has been initialized and you do not need to initialize it again. Initializing the switch resets the entire internal Flash device and erases all data in the bootflash: file system. The internal Flash device is composed of several file systems with bootflash: being one of them. All files in bootflash: are erased and you must download the system and kickstart images again. After initializing the switch, you do not have to format the bootflash: again because it is automatically formatted.

If the bootflash: is found corrupted during a boot sequence, you will see the following message on the CLI:

```
ERROR:bootflash: has unrecoverable error; please do "format bootflash:"
```

Refer to the *Cisco MDS 9000 Family CLI Configuration Guide* for more information on the format bootflash command.

To initialize the bootflash using Device Manager, follow these steps:

Step 1 Click **Admin** > **Reset Switch**.

Step 2 Click **Reset System**.

Using the File System

All switches in the Cisco MDS 9000 Family have one internal bootflash: that resides in the supervisor or switching module.You have access to two directories within the internal bootflash: file system.

- The volatile: directory provides temporary storage, and it is also the default. Files in temporary storage (volatile:) are erased when the switch reboots.

- The bootflash: (nonvolatile storage) directory provides permanent storage. Files in permanent storage (bootflash:) are preserved through reboots and power outages.

Cisco MDS 9500 Series directors contain an additional external CompactFlash referred to as the slot0: directory. The external CompactFlash, an optional device for MDS 9500 Series directors, can be used for storing software images, logs, and core dumps.

You can use Device Manager to perform the following functions to help you manage software image files and configuration files:

- Figure 14-4Flash Files, page 14-4
- Creating a Directory, page 14-3
- Deleting an Existing File or Directory, page 14-4
- Copying Files, page 14-5
- Performing Other File Manipulation Tasks, page 14-7

To list the files in a directory using Device Manager, follow these steps:

Step 1 Click **Admin > Flash File**s.

By default, you see the bootflash directory listed for the supervisor's local partition.

Figure 14-3 Flash Files

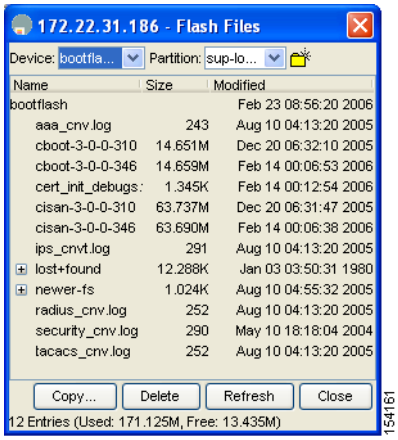

Step 2 Select the device and partition for the directory you want to view from the drop-down lists.

You see a list of files and directories.

Creating a Directory

To create a directory using Device Manager, follow these steps:

Step 1 Click **Admin > Flash Files**.

By default, you see the bootflash directory listed for the supervisor's local partition.

Figure 14-4 Flash Files

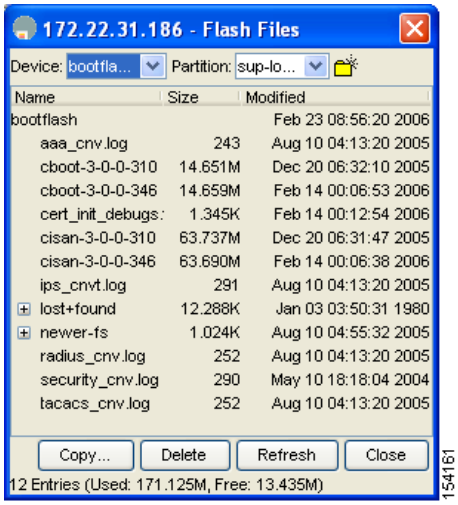

Step 2 Select the device and partition from the drop-down lists for the directory where you want to create the directory.

Step 3 Click the **Create Directory** icon (yellow folder in Figure 14-4) to create a directory.

You see the Create New Directory dialog box shown in Figure 14-5.

Figure 14-5 New Directory Dialog Box

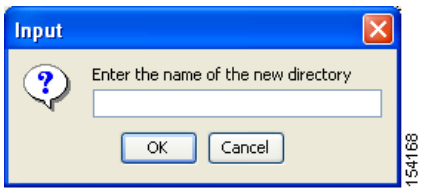

Step 4 Enter the name of the new directory, and click **OK** (see Figure 14-5).

You see the new directory in the directory listing.

Tip Any directory saved in the volatile: file system is erased when the switch reboots.

Deleting an Existing File or Directory

To delete a file or directory using Device Manager, follow these steps:

Step 1 Click **Admin > Flash Files**.

By default, you see the bootflash: directory listed for the supervisor's local partition shown in Figure 14-6.

Figure 14-6 Flash Files

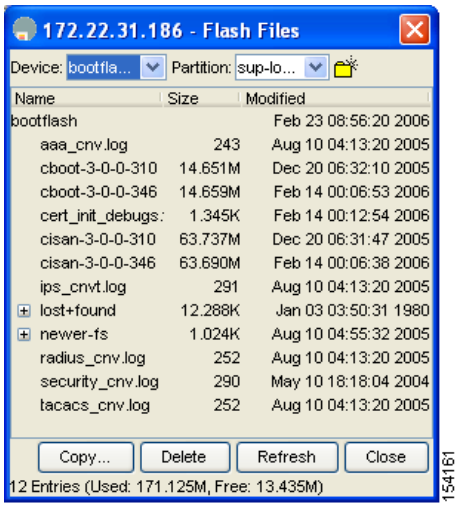

Step 2 Select a device and partition from the drop-down lists.

Step 3 Click the file or directory you want to delete.

Step 4 Click **Delete** to delete the file or directory.

Caution If you specify a directory, the delete removes the entire directory and all of its contents.

Copying Files

To copy a file using Device Manager, follow these steps:

Step 1 Select **Admin > Flash Files**.

By default, you see the bootflash: directory listed for the supervisor's local partition shown in Figure 14-7.

Figure 14-7 Flash Files

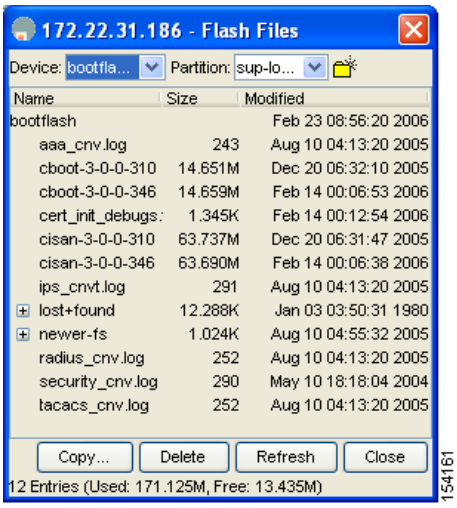

Step 2 Select the device and partition from the drop-down lists for the directory containing the file you want to copy.

Step 3 Click **Copy.**

You see the Copy dialog box.

Figure 14-8 Copy FLash Files in Device Manager

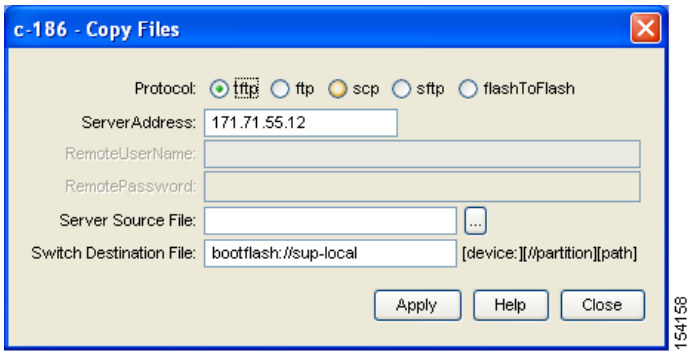

Step 4 Choose the protocol you want to use for the copy, **tftp**, **ftp**, **scp**, or **flashToFlash.**

Step 5 Enter the address of the source server for a Flash to Flash copy only.

Step 6 Click the **...** button to browse for the source file on your local PC or on the server, depending on the type of copy.

Note If you are copying from Flash, the file name must be in the format: [*device*>:][<*partition*>:]<*file*>

where *device* is a value obtained from FlashDeviceName, *partition* is a value obtained from FlashPartitionName and *file* is the name of a file in Flash.

Step 7 Enter the Switch Destination File name. (See the note in Step 6.)

Step 8 Click **Apply.**

Performing Other File Manipulation Tasks

To perform the following CLI-specific tasks, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*:

- Displaying file contents
- Displaying the last line in a file
- Saving output to a file
- Moving files
- Compressing and uncompressing files
- Executing commands specified in a script
- Setting the delay time

Working with Configuration Files

Configuration files can contain some or all of the commands needed to configure one or more switches. For example, you might want to download the same configuration file to several switches that have the same hardware configuration so that they have identical module and port configurations.

This section describes how to work with configuration files and has the following topics:

Downloading Configuration Files to the Switch

You can configure a switch in the Cisco MDS 9000 Family by using configuration files you create or download from another switch. In addition, you can store configuration files on a bootflash device on the supervisor module and you can configure the switch using a configuration stored on an external CompactFlash disk.

Before you begin downloading a configuration file using a remote server, do the following:

- Ensure the configuration file to be downloaded is in the correct directory on the remote server.
- Ensure that the permissions on the file are set correctly. Permissions on the file should be set to world-read.
- Ensure the switch has a route to the remote server. The switch and the remote server must be in the same subnetwork if you do not have a router or default gateway to route traffic between subnets.

  Check connectivity to the remote server using **ping**.

Saving the Configuration

To save the configuration file using Device Manager, follow these steps:

Step 1 Click **Admin > Save Configuration**.

You see the message `Really save running to startup configuration?`

Step 2 Click **Yes** to save the configuration. Click **No** to close the pop-up window without saving the configuration.

After you have created a running configuration in system memory, you can save it to the startup configuration in NVRAM.

To save the configuration file using Device Manager, follow these steps:

Step 1 Click **Admin > Copy Configuration**.

You see the Copy Configuration dialog box shown in Figure 14-9.

Figure 14-9 Copy Configuration Dialog Box

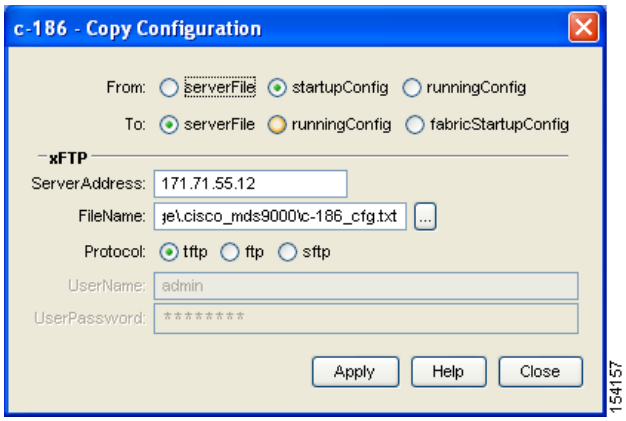

Step 2 Select the location of the file you will copy from (**serverFile**, **startupConfig**, **runningConfig**).

Step 3 Select the location of the file you will copy to (**serverFile**, **runningConfig**, **fabricStartupConfig**).

Note You can copy a file fabric-wide using the fabricStartupConfig option, available in Cisco MDS SAN-OS Release 2.1(1a) or later.

Step 4 Enter the server address of the source server.

Step 5 Click the **...** button to browse for the source file on the switch or the server, depending on the type of copy.

Step 6 Select the protocol you want to use to perform the copy procedure, **tftp**, **ftp**, or **sftp.**

Step 7 Enter the user name and password you use to access the switch or server.

Step 8 Click **Apply.**

Saving Startup Configurations in the Fabric

You can use Cisco Fabric Services (CFS) to instruct the other switches in the fabric to save their configurations to their local NVRAM. You can copy the running configuration to the startup configuration across the entire fabric by using the fabricStartupConfig option. This triggers every switch in the fabric to copy its running configuration to its startup configuration.

Note If any switch fails during this fabric-wide copy, that switch and the switch that you used to initiate this process will keep the existing startup configuration. This does not affect the other switches in the fabric.

Backing Up the Current Configuration

Before installing or migrating to any software configuration, back up the startup configuration. To back up the startup configuration using Device Manager, follow these steps:

Step 1 Select **Admin > Copy Configuration**.

You see the Copy Configuration dialog box shown in Figure 14-10.

Figure 14-10 Copy Configuration Dialog Box

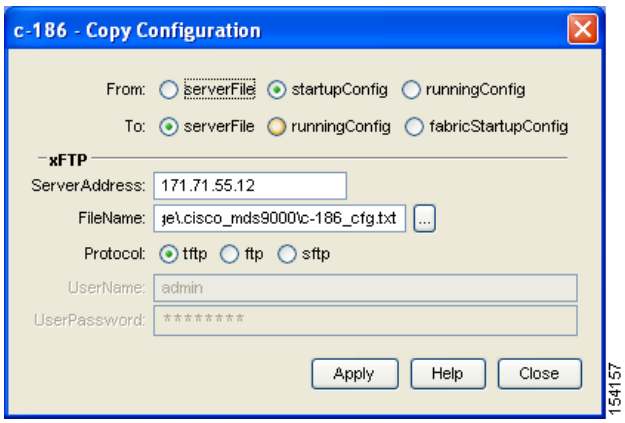

Step 2 Select the location of the file you want to back up (**server file**, **startup configuration**, or **running configuration**).

Step 3 Select the destination of the file (**server file**, **running configuration**, **fabric startup configuration**).

Step 4 Enter the server address.

Step 5 Click the **...** button to select the file name.

Step 6 Choose the file transfer protocol (**tftp**, **ftp**, or **sftp**).

Step 7 Enter the user name and password for the server you specified in Step 4.

Step 8 Click **Apply** to copy the file.

CHAPTER 15

Configuring High Availability

The Cisco MDS 9500 Series of multilayer directors support application restartability and nondisruptive supervisor switchability. The switches are protected from system failure by redundant hardware components and a high availability software framework.

This chapter includes the following sections:

About High Availability

The high availability (HA) software framework provides the following:

- Ensures nondisruptive software upgrade capability. See Chapter 13, “Software Images.”
- Provides redundancy for supervisor module failure by using dual supervisor modules.
- Performs nondisruptive restarts of a failed process on the same supervisor module. A service running on the supervisor modules and on the switching module tracks the HA policy defined in the configuration and takes action based on this policy. This feature is also available in switches in the Cisco MDS 9200 Series and the Cisco MDS 9100 Series.
- Protects against link failure using the PortChannel (port aggregation) feature. This feature is also available in switches in the Cisco MDS 9200 Series and in the Cisco MDS 9100 Series. See Chapter 21, “Configuring PortChannels.”
- Provides management redundancy using the Virtual Router Redundancy Protocol (VRRP). This feature is also available in switches in the Cisco MDS 9200 Series and in the Cisco MDS 9100 Series.

  See the “Virtual Router Redundancy Protocol” section on page 46-8.
- Provides switchovers if the active supervisor fails, the standby supervisor, if present, takes over without disrupting storage or host traffic.

  Directors in the Cisco MDS 9500 Series have two supervisor modules (sup-1 and sup-2) in slots 5 and 6 (Cisco MDS 9509 and 9506 Switches) or slots 7 and 8 (Cisco MDS 9513 Switch). When the switch powers up and both supervisor modules are present, the supervisor module that comes up first enters the active mode and the supervisor module that comes up second enters the standby mode. If

both supervisor modules come up at the same time, sup-1 becomes active. The standby supervisor module constantly monitors the active supervisor module. If the active supervisor module fails, the standby supervisor module takes over without any impact to user traffic.

Switchover Mechanisms

Switchovers occur by one of the following two mechanisms:

- The active supervisor module fails and the standby supervisor module automatically takes over.
- You manually initiate a switchover from an active supervisor module to a standby supervisor module.

Once a switchover process has started another switchover process cannot be started on the same switch until a stable standby supervisor module is available.

Note If the standby supervisor module is not in a stable state (ha-standby), a switchover is not performed.

HA Switchover Characteristics

An HA switchover has the following characteristics:

- It is stateful (nondisruptive) because control traffic is not impacted.
- It does not impact data traffic because the switching modules are not impacted.
- Switching modules are not reset.

Initiating a Switchover

To manually initiate a switchover from an active supervisor module to a standby supervisor module, reset the active supervisor module using Device Manager. Once the switchover process has started, another switchover process cannot be started on the same switch until a stable standby module is available.

To perform a switchover using Device Manager, follow these steps:

Step 1 Ensure that an HA switchover is possible by clicking **Physical > Modules** to verify the presence of multiple modules. See the example in Figure 15-1.

Figure 15-1 Modules Screen Shows Current Supervisor

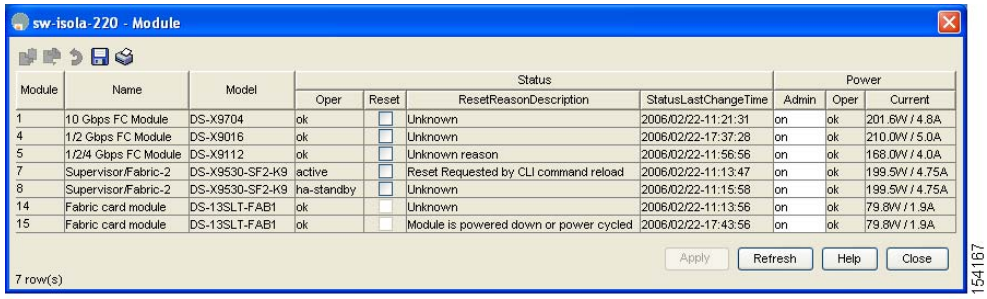

Step 2 Click **Admin** > **Reset Switch** on the main Device Manager screen.

Figure 15-2 Reset Switch Dialog Box

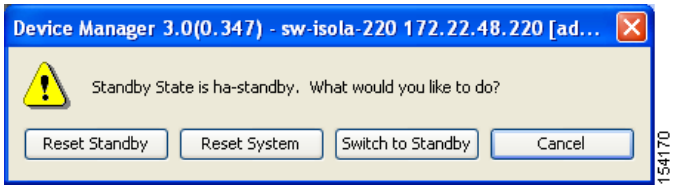

Step 3 Click **Switch to Standby**.

Switchover Guidelines

Be aware of the following guidelines when performing a switchover:

- When you manually initiate a switchover, system messages indicate the presence of two supervisor modules.
- A switchover can only be performed when two supervisor modules are functioning in the switch.
- The modules in the chassis are functioning as designed.

Process Restartability

Process restartability provides the high availability functionality in Cisco MDS 9000 Family switches. It ensures that process-level failures do not cause system-level failures. It also restarts the failed processes automatically. This vital process functions on infrastructure that is internal to the switch.

See the "Displaying System Processes" section on page 64-1.

Synchronizing Supervisor Modules

The running image is automatically synchronized in the standby supervisor module by the active supervisor module. The boot variables are synchronized during this process.

The standby supervisor module automatically synchronizes its image with the running image on the active supervisor module.

See the "Replacing Modules" section on page 13-16.

The following conditions identify when automatic synchronization is possible:

- If the internal state of one supervisor module is `Active with HA standby` and the other supervisor module is `HA-standby`, the switch is operationally HA and can do automatic synchronization.
- If the internal state of one of the supervisor modules is `none`, the switch cannot do automatic synchronization.

Table 15-1 lists the possible values for the redundancy states.

Table 15-1 Redundancy States

State	Description
Not present	The supervisor module is not present or is not plugged into the chassis.
Initializing	The diagnostics have passed and the configuration is being downloaded.
Active	The active supervisor module and the switch is ready to be configured.
Standby	A switchover is possible.
Failed	The switch detects a supervisor module failure on initialization and automatically attempts to power-cycle the module three (3) times. After the third attempt it continues to display a failed state.
Offline	The supervisor module is intentionally shut down for debugging purposes.
At BIOS	The switch has established connection with the supervisor and the supervisor module is performing diagnostics.
Unknown	The switch is in an invalid state. If it persists, call TAC.

Table 15-2 lists the possible values for the supervisor module states.

Table 15-2 Supervisor States

State	Description
Active	The active supervisor module in the switch is ready to be configured.
HA standby	A switchover is possible.
Offline	The switch is intentionally shut down for debugging purposes.
Unknown	The switch is in an invalid state and requires a support call to TAC.

Table 15-3 lists the possible values for the internal redundancy states.

Table 15-3 Internal States

State	Description
HA standby	The HA switchover mechanism in the standby supervisor module is enabled (see the "HA Switchover Characteristics" section on page 15-2).
Active with no standby	A switchover is possible.
Active with HA standby	The active supervisor module in the switch is ready to be configured. The standby module is in the HA-standby state.
Shutting down	The switch is being shut down.
HA switchover in progress	The switch is in the process of changing over to the HA switchover mechanism.
Offline	The switch is intentionally shut down for debugging purposes.
HA synchronization in progress	The standby supervisor module is in the process of synchronizing its state with the active supervisor modules.
Standby (failed)	The standby supervisor module is not functioning.
Active with failed standby	The active supervisor module and the second supervisor module is present but is not functioning.
Other	The switch is in a transient state. If it persists, call TAC.

CHAPTER 16

Managing System Hardware

This chapter provides details on how to manage system hardware other than services and switching modules and how to monitor the health of the switch. It includes the following sections:

Displaying Switch Hardware Inventory

To view information on the field replaceable units (FRUs) in the switch, including product IDs and serial numbers, follow these steps:

Step 1 In Fabric Manager , choose a fabric or switch in the Logical Domains pane, then expand **Switches** and select **Hardware** in the Physical Attributes pane.

You see a list like the one shown in Figure 16-1.

Figure 16-1 Fabric Manager Hardware Inventory

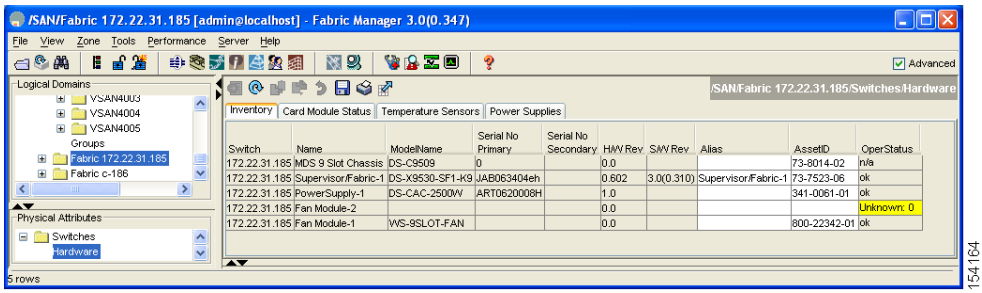

In Device Manager, click **Physical > Inventory**. You see a list like the one shown in Figure 16-2.

Figure 16-2 Device Manager Hardware Inventory

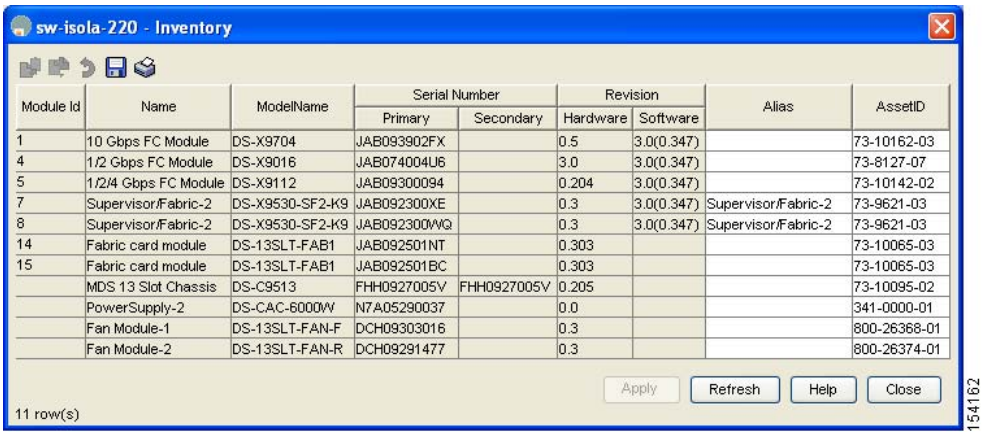

You see system attributes for multiple switches in Figure 16-1 and Figure 16-2. To see attributes for a single switch in Device Manager, double click the graphic of the switch in the main screen.

Note To configure modules, see Chapter 17, "Managing Modules."

Displaying the Switch Serial Number

The serial number of your Cisco MDS 9000 Family switch can be obtained by looking at the serial number label on the back of the switch (next to the power supply) or from Fabric Manager by selecting that switch in the Logical Domains pane, then expanding **Switches** and selecting **Hardware** in the Physical Attributes pane in Fabric Manager. The Serial No Primary column in the Information pane shows the serial number.

Displaying Power Usage Information

Use Fabric Manager to display power usage. Select a switch in the Logical Domains pane, expand **Switches** and select **Hardware** in the Physical Attributes pane, then click the **Power Supplies** tab in the Information pane to display actual power usage information for the entire switch. See the first example under Power Supply Configuration Modes.

Note In a Cisco MDS 9500 Series switch, power usage is reserved for both supervisors whether one or both supervisor modules are present.

Power Supply Configuration Modes

Switches in the MDS 9000 Family have two redundant power supply slots. The power supplies can be configured in either redundant or combined mode.

- Redundant mode—Uses the capacity of one power supply only. This is the default mode. In case of power supply failure, the entire switch has sufficient power available in the system.
- Combined mode—Uses the combined capacity of both power supplies. In case of power supply failure, the entire switch can be shut down (depends on the power used) causing traffic disruption. This mode is seldom used, except in cases where the switch has two low power supply capacities but a higher power usage.

Note The chassis in the Cisco MDS 9000 Family uses 1200 W when powered at 110 V, and 2500 W when powered at 220 V.

To configure the power supply mode, follow these steps:

Step 1 In the Fabric Manager Physical Attributes pane, expand **Switches** and then select **Hardware.** Click the **Power Supplies** tab.

You see the power supply information screen shown in Figure 16-3.

Figure 16-3 Power Supply Information in Fabric Manager

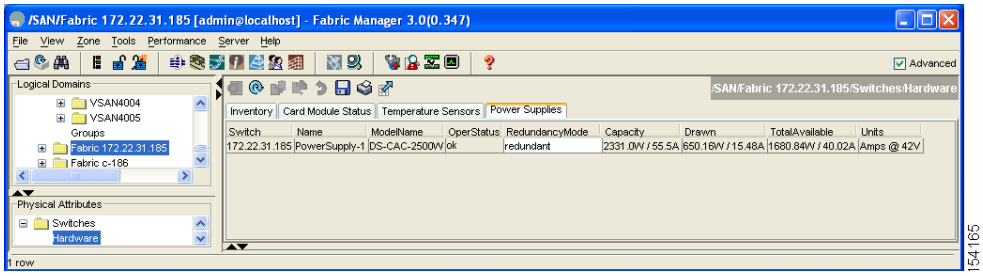

In Device Manager, click **Physical > Power Supplies**.

You see the screen in Figure 16-4.

Figure 16-4 Power Supply Information in Device Manager

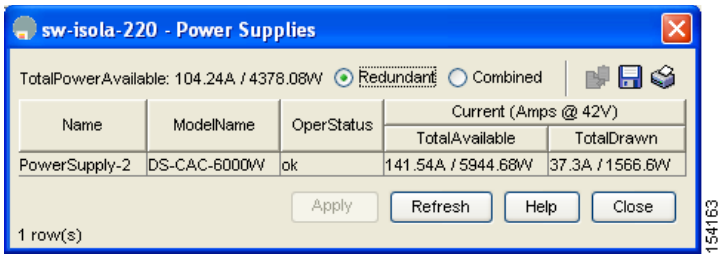

Step 2 Configure the power attributes for the power supply.

Step 3 Click **Apply** in Device Manager or click the **Apply Changes** icon in Fabric Manager.

Note See the "Displaying Power Usage Information" section on page 16-3 to view the current power supply configuration.

Power Supply Configuration Guidelines

Follow these guidelines when configuring power supplies:

1. When power supplies with different capacities are installed in the switch, the total power available differs based on the configured mode, either redundant or combined:

   a. Redundant mode—the total power is the lesser of the two power supply capacities.
   For example, suppose you have the following usage figures configured:

   Power supply 1 = 2500 W
   Additional power supply 2 = not used
   Current usage = 2000 W
   Current capacity = 2500 W

   Then the following three scenarios differ as specified (see Table 16-1):

   Scenario 1: If 1800 W is added as power supply 2, then power supply 2 is shut down.
   Reason: 1800 W is less than the usage of 2000 W.

   Scenario 2: If 2200 W is added as power supply 2, then the current capacity decreases to 2200 W.
   Reason: 2200 W is the lesser of the two power supplies.

   Scenario 3: If 3000 W is added as power supply 2, then the current capacity value remains at 2500 W.
   Reason: 2500 W is the lesser of the two power supplies.

Table 16-1 Redundant Mode Power Supply Scenarios

Scenario	Power Supply 1 (W)[1]	Current Usage (W)	Insertion of Power Supply 2 (W)	New Capacity (W)	Action Taken by Switch
1	2500	2000	1800	2500	Power supply 2 is shut down.

Table 16-1 Redundant Mode Power Supply Scenarios (continued)

Scenario	Power Supply 1 (W)[1]	Current Usage (W)	Insertion of Power Supply 2 (W)	New Capacity (W)	Action Taken by Switch
2	2500	2000	2200	2200	Capacity becomes 2200 W.
3	2500	2000	3300	2500	Capacity remains the same.

1. W = Watts

b. Combined mode—the total power is twice the lesser of the two power supply capacities.

For example, suppose you have the following usage figures configured:

Power supply 1 = 2500 W
Additional Power supply 2 = not used
Current Usage = 2000 W
Current capacity = 2500 W

Then, the following three scenarios differ as specified (see Table 16-2):

Scenario 1: If 1800 W is added as power supply 2, then the capacity increases to 3600 W.
Reason: 3600 W is twice the minimum (1800 W).

Scenario 2: If 2200 W is added as power supply 2, then the current capacity increases to 4400 W.
Reason:4400 W is twice the minimum (2200 W).

Scenario 3: If 3000 W is added as power supply 2, then the current capacity increases to 5000 W.
Reason: 5000 W is twice the minimum (2500 W).

Table 16-2 Combined Mode Power Supply Scenarios

Scenario	Power Supply 1 (W)[1]	Current Usage (W)	Insertion of Power Supply 2 (W)	New Capacity (W)	Action Taken by Switch
1	2500	2000	1800	3600	Power is never shut down. The new capacity is changed.
2	2500	2000	2200	4400	
3	2500	2000	3300	5000	

1. W = Watts

2. When you change the configuration from combined to redundant mode and the system detects a power supply that has a capacity lower than the current usage, the power supply is shut down. If both power supplies have a lower capacity than the current system usage, the configuration is not allowed. Several configuration scenarios are summarized in Table 16-3.

Scenario 1: You have the following usage figures configured:

Power supply 1 = 2500 W
Additional Power supply 2 = 1800 W
Current Usage = 2000 W
Current mode = combined mode (so current capacity is 3600 W)

You decide to change the switch to redundant mode. Then power supply 2 is shut down.

Reason: 1800 W is the lesser of the two power supplies and it is less than the system usage.

Scenario 2: You have the following usage figures configured:

Power supply 1 = 2500 W
Additional Power supply 2 = 2200 W
Current Usage = 2000 W
Current mode = combined mode (so current capacity is 4400 W).

You decide to change the switch to redundant mode. Then the current capacity decreases to 2200 W.

Reason: 2200 W is the lesser of the two power supplies.

Scenario 3: You have the following usage figures configured:

Power supply 1 = 2500 W
Additional Power supply 2 = 1800 W
Current Usage = 3000 W
Current mode = combined mode (so current capacity is 3600 W).

You decide to change the switch to redundant mode. Then the current capacity decreases to 2500 W and the configuration is rejected.

Reason: 2500 W is less than the system usage (3000 W).

Table 16-3 Combined Mode Power Supply Scenarios

Scenario	Power Supply 1 (W)[1]	Current Mode	Current Usage (W)	Power Supply 2 (W)	New Mode	New Capacity (W)	Action Taken by Switch
1	2500	combined	2000	1800	N/A	3600	This is the existing configuration.
	2500	N/A	2000	1800	redundant	2500	Power supply 2 is shut down
2	2500	combined	2000	2200	N/A	4400	This is the existing configuration.
	2500	N/A	2000	2200	redundant	2200	The new capacity is changed.
3	2500	combined	3000	1800	N/A	3600	This is the existing configuration.
	2500	N/A	3000	1800	redundant	N/A	Rejected, so the mode reverts to combined mode.

1. W = Watts

About Crossbar Management

Cisco MDS SAN-OS Release 3.0(1) and later supports two types of hardware for the Cisco MDS 9500Series Directors: Generation 1 and Generation 2.

Generation 1 consists of all hardware supported by Cisco SAN-OS prior to Release 3.0(1), including the following:

- Cisco MDS 9506 and 9509 Director chassis
- Supervisor-1 module
- 32-port 2-Gbps Fibre Channel switching module

- 16-port 2-Gbps Fibre Channel switching module
- 8-port IP Storage Services (IPS-8) module
- 4-port IP Storage Services (IPS-4) module
- Storage Services Module (SSM)
- 14/2 port Multiprotocol Services (MPS-14/2) module

Generation 2 consists of all new hardware support by Cisco SAN-OS Release 3.0(1) and later, including the following:

- Cisco MDS 9513 Director chassis
- Supervisor-2 module
- 48-port 4-Gbps Fibre Channel switching module
- 24-port 4-Gbps Fibre Channel switching module
- 12-port 4-Gbps Fibre Channel switching module
- 4-port 10-Gbps Fibre Channel switching module

The Cisco MDS 9500 Series Directors running Cisco MDS SAN-OS 3.0(1) or later support the following types of crossbars:

- Integrated crossbar—Located on the Supervisor 1 and Supervisor 2 modules. The Cisco MDS 9506 and 9509 Directors only use integrated crossbars.
- External crossbar—Located on an external crossbar switching module. External crossbar switching modules are required for Cisco MDS 9513 Directors.

Operational Considerations When Removing Crossbars

You can mix and match Generation 1 and Generation 2 hardware on the Cisco MDS 9500 Series Directors running Cisco MDS SAN-OS 3.0(1) or later without compromising the integrity and availability of your SANs based on Cisco MDS 9500 Series Directors.

To realize these benefits, you must consider the following operational requirements when removing crossbars for maintenance activities:

- Graceful Shutdown of a Crossbar, page 16-7
- Backward Compatibility for Generation 1 Modules in Cisco MDS 9513 Directors, page 16-9

Graceful Shutdown of a Crossbar

You must perform a graceful shutdown of a crossbar (integrated or external) before removing it from the MDS 9500 Series Director.

To perform a graceful shutdown of external crossbar switching modules in a Cisco MDS 9513 Director using Device Manager, follow these steps:

Step 1 Right-click the supervisor module. Crossbars are indicated with a green X (see Figure 16-5).

You see the context menu for the supervisor module.

Figure 16-5 Shutting Down a Crossbar

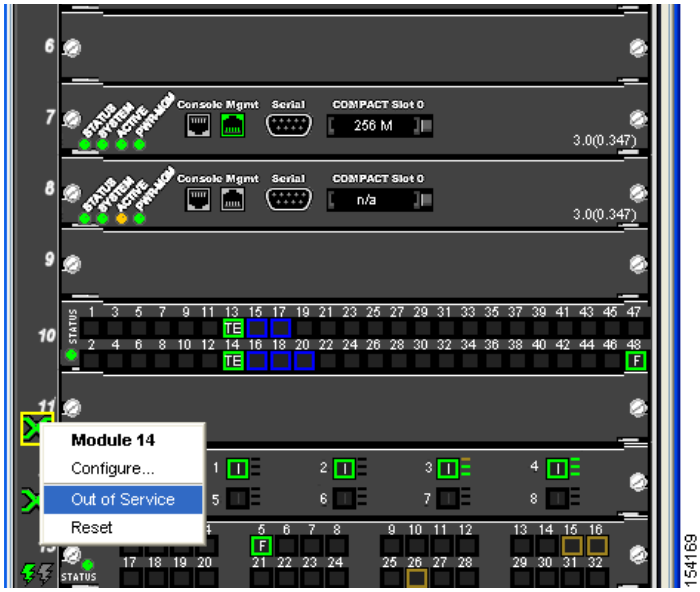

Step 2 Select **Out of Service** to gracefully shut down the integrated crossbar.

Note To reactivate the integrated crossbar, you must remove and reinsert or replace the Supervisor-1 or Supervisor-2 module.

Caution Taking the crossbar out of service may cause a supervisor switchover.

To perform a graceful shutdown of integrated crossbars on the supervisor module in a Cisco MDS 9509 or 9506 Director using Device Manager, follow these steps::

Step 1 Right-click the external crossbar switching module.

You see the context menu for that module.

Step 2 Select **Out of Service** to gracefully shut down the external crossbar switching module.

Note To reactivate the external crossbar module, you must remove and reinsert or replace the crossbar module.

Caution Taking the crossbar out of service may cause a supervisor switchover.

Backward Compatibility for Generation 1 Modules in Cisco MDS 9513 Directors

To provide backward compatibility for a Generation 1 module in a Cisco MDS 9513 chassis, the active and backup Supervisor-2 modules are associated to a specific crossbar module. The Supervisor-2 module in slot 7 is associated with crossbar module 1 and Supervisor-2 module in slot 8 is associated with crossbar module 2. You must plan for the following operational considerations before removing crossbar modules:

- Whenever a crossbar module associated with the active Supervisor-2 module goes offline or is brought online in a system that is already online, a stateful supervisor switchover occurs. This switchover does not disrupt traffic. Events that cause a crossbar module to go offline include the following:Out-of-service requests
- Physical removal
- Errors
- Supervisor-2 module switchovers do not occur if the crossbar switching module associated with the backup Supervisor-2 module goes offline.

Note Supervisor-2 module switchovers do not occur when removing crossbar switch modules on a Cisco MDS 9513 that only has Generation-2 modules installed.

About Module Temperature

Built-in, automatic sensors are provided in all switches in the Cisco MDS 9000 Family to monitor your switch at all times.

Each module (switching and supervisor) has four sensors: 1 (outlet sensor), 2 (intake sensor), 3 (onboard sensor), and 4 (onboard sensor). Each sensor has two thresholds (in degrees Celsius): minor and major.

Note A threshold value of -127 indicates that no thresholds are configured or applicable.

- Minor threshold—When a minor threshold is exceeded, a minor alarm occurs and the following action is taken for all four sensors:
  - System messages are displayed.
  - Call Home alerts are sent (if configured).
  - SNMP notifications are sent (if configured).
- Major threshold—When a major threshold is exceeded, a major alarm occurs and the following action is taken.
  - For sensors 1, 3, and 4 (outlet and onboard sensors):

    System messages are displayed.

    Call Home alerts are sent (if configured).

    SNMP notifications are sent (if configured).
  - For sensor 2 (intake sensor):

    If the threshold is exceeded in a switching module, only that module is shut down.

    If the threshold is exceeded in an active supervisor module with HA-standby or standby present, only that supervisor module is shut down and the standby supervisor module takes over.

    If you do not have a standby supervisor module in your switch, you have an interval of 2 minutes to decrease the temperature. During this interval the software monitors the temperature every five (5) seconds and continuously sends system messages as configured.

Tip To realize the benefits of these built-in, automatic sensors on any switch in the Cisco MDS 9500 Series, we highly recommend that you install dual supervisor modules. If you are using a Cisco MDS 9000 Family switch without dual supervisor modules, we recommend that you immediately replace the fan module if even one fan is not working.

Displaying Module Temperature

Expand **Switches** and then select **Hardware** in the Physical Attributes pane in Fabric Manager. Click the **Temperature Sensors** tab in the Information pane to display temperature sensors for each module (see the second example under Power Supply Configuration Modes).

About Fan Modules

Hot-swappable fan modules (fan trays) are provided in all switches in the Cisco MDS 9000 Family to manage airflow and cooling for the entire switch. Each fan module contains multiple fans to provide redundancy. The switch can continue functioning in the following situations:

- One or more fans fail within a fan module—Even with multiple fan failures, switches in the Cisco MDS 9000 Family can continue functioning. When a fan fails within a module, the functioning fans in the module increase their speed to compensate for the failed fan(s).
- The fan module is removed for replacement—The fan module is designed to be removed and replaced while the system is operating without presenting an electrical hazard or damage to the system. When replacing a failed fan module in a running switch, be sure to replace the new fan module within five minutes.

Tip If one or more fans fail within a fan module, the Fan Status LED turns red. A fan failure could lead to temperature alarms if not corrected immediately.

The fan status is continuously monitored by the Cisco MDS SAN-OS software. In case of a fan failure, the following action is taken:

- System messages are displayed.
- Call Home alerts are sent (if configured).
- SNMP notifications are sent (if configured).

To display the fan module status, from Device Manager, choose **Physical > Fans**. The dialog box displays the fan status.

The possible Status field values for a fan module on the Cisco MDS 9500 Series switches are as follows:

- If the fan module is operating properly, the status is ok.
- If the fan is physically absent, the status is absent.
- If the fan is physically present but not working properly, the status is failure.

On the Cisco MDS 9513 Director, the front fan module has 15 fans.

Figure 16-6 shows the numbering of the fans in the front fan module on the Cisco MDS 9513 Director.

Figure 16-6 Cisco MDS 9513 Front Fan Module Numbering

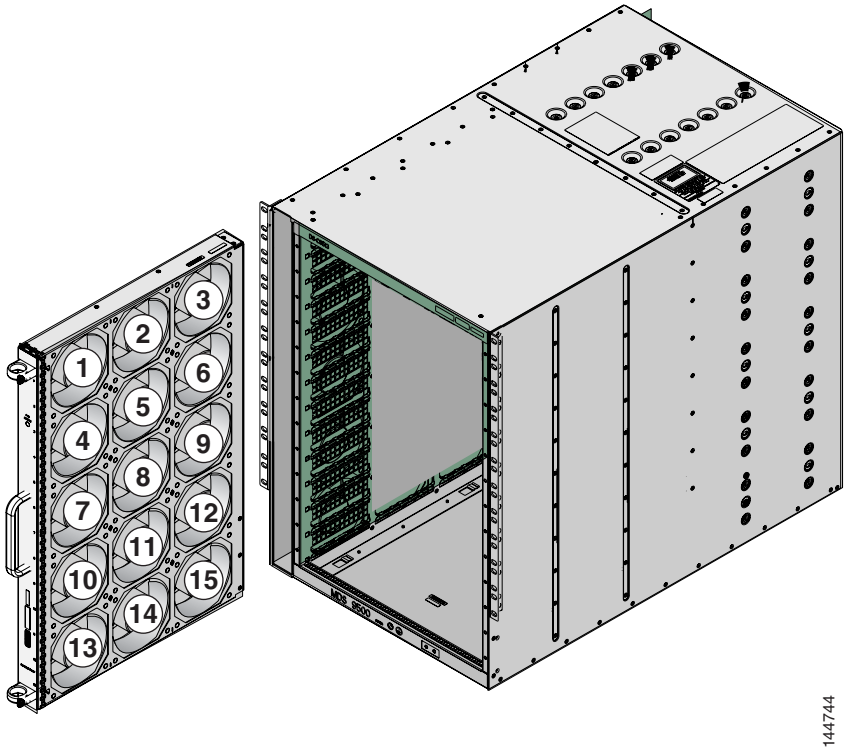

The rear fan module (DS-13SLT-FAN-R) on the Cisco MDS 9513 Director has only two fans.

Default Settings

Table 16-4 lists the default hardware settings.

Table 16-4 Default Hardware Parameters

Parameters	Default
Power supply mode	Redundant mode.

CHAPTER 17

Managing Modules

This chapter describes how to manage switching and services modules (also known as line cards) and provides information on monitoring module states. This chapter includes the following sections:

About Modules

Table 17-1 describes the supervisor module options for switches in the Cisco MDS 9000 Family.

Table 17-1 Supervisor Module Options

Product	Number of Supervisor Modules	Supervisor Module Slot Number	Switching and Services Module Features
Cisco MDS 9513	Two modules	7 and 8	13-slot chassis allows any switching or services module in the other eleven slots.
Cisco MDS 9509	Two modules	5 and 6	9-slot chassis allows any switching or services module in the other seven slots.
Cisco MDS 9506	Two modules	5 and 6	6-slot chassis allows any switching or services module in the other four slots.
Cisco MDS 9216	One module	1	2-slot chassis allows one optional switching or services module in the other slot.

Table 17-1 Supervisor Module Options (continued)

Product	Number of Supervisor Modules	Supervisor Module Slot Number	Switching and Services Module Features
Cisco MDS 9216A	One module	1	2-slot chassis allows one optional switching or services module in the other slot.
Cisco MDS 9216i	One module	1	2-slot chassis allows one optional switching or services module in the other slot.

Supervisor Modules

Supervisor modules are automatically powered up and started with the switch.

- Cisco MDS 9513 Directors have two supervisor modules—one in slot 7 (sup-1) and one in slot 8 (sup-2). See Table 17-2. When the switch powers up and both supervisor modules come up together, the active module is the one that comes up first. The standby module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.
- Cisco MDS 9506 and Cisco MDS 9509 switches have two supervisor modules—one in slot 5 (sup-1) and one in slot 6 (sup-2). See Table 17-2. When the switch powers up and both supervisor modules come up together, the active module is the one that comes up first. The standby module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.
- Cisco MDS 9216i switches have one supervisor module that includes an integrated switching module with 14 Fibre Channel ports and two Gigabit Ethernet ports.
- Cisco MDS 9200 Series switches have one supervisor module that includes an integrated 16-port switching module.

Table 17-2 Supervisor Module Terms and Usage

Module Terms	Fixed or Relative	Usage
`module-7 and module-8`	Fixed usage for MDS 9513	module-7 always refers to the supervisor module in slot 7 and module-8 always refers to the supervisor module in slot 8.
`module-5 and module-6`	Fixed usage for MDS 9509 and MDS 9506	module-5 always refers to the supervisor module in slot 5 and module-6 always refers to the supervisor module in slot 6.
`module-1`	Fixed usage for MDS 9200 series	module-1 always refers to the supervisor module in slot 1.
`sup-1 and sup-2`	Fixed usage	On the MDS 9506 and MDS 9509 switches, sup-1 always refers to the supervisor module in slot 5 and sup-2 always refers to the supervisor module in slot 6. On the MDS 9513 Directors, sup-1 always refers to the supervisor module in slot 7 and sup-2 always refers to the supervisor module in slot 8.

Table 17-2 Supervisor Module Terms and Usage (continued)

Module Terms	Fixed or Relative	Usage
`sup-active and sup-standby`	Relative usage	sup-active refers to the active supervisor module—relative to the slot that contains the active supervisor module. sup-standby refers to the standby supervisor module—relative to the slot that contains the standby supervisor module.
`sup-local and sup-remote`	Relative usage	If you are logged into the active supervisor, sup-local refers to the active supervisor module and sup-remote refers to the standby supervisor module. If you are logged into the standby supervisor, sup-local refers to the standby supervisor module (the one you are logged into.) There is no sup-remote available from the standby supervisor module (you cannot access a file system on the active sup).

Switching Modules

Cisco MDS 9000 Family switches support any switching module in any non-supervisor slot. These modules obtain their image from the supervisor module.

Services Modules

Cisco MDS 9000 Family switches support any services module in any non-supervisor slot.

Refer to the *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for more information on CSMs.

Verifying the Status of a Module

Before you begin configuring the switch, you need to ensure that the modules in the chassis are functioning as designed. To verify the status of a module at any time expand **Switches** and then select **Hardware** in the Physical Attributes pane in Fabric Manager and select **Card Module Status** tab in the Information pane (see the "Fibre Channel Interfaces" section on page 18-1). The interfaces in each module are ready to be configured when the **ok** status is displayed. A sample screenshot follows:

Figure 17-1 Card Module Status Display

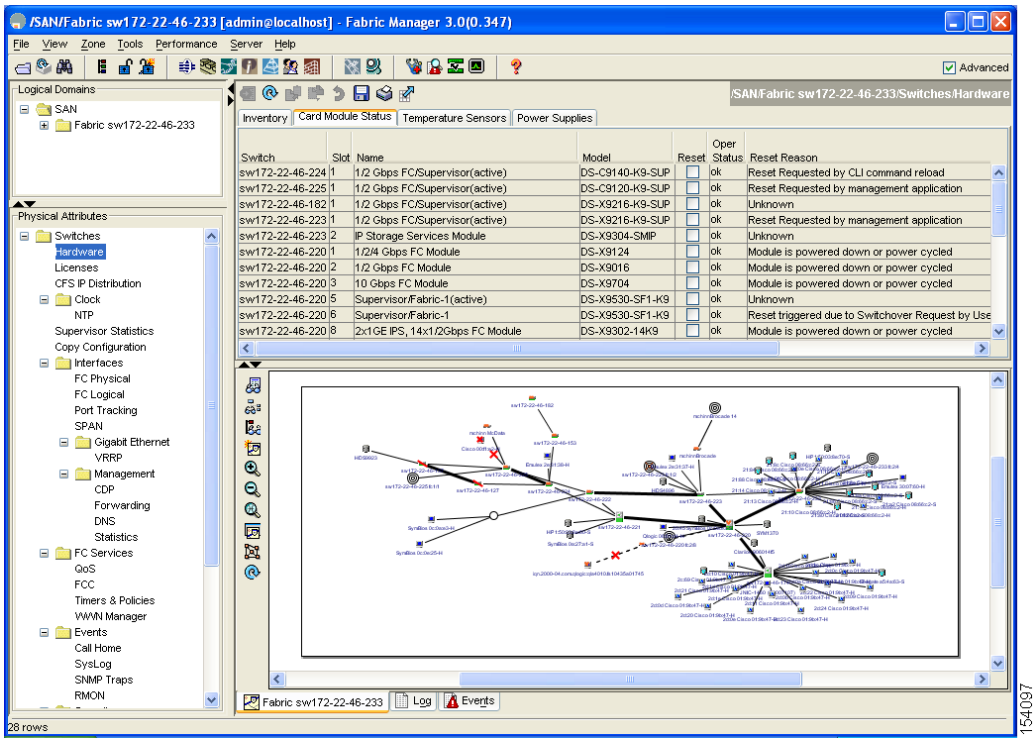

The Status column in the output should display an **ok** status for switching modules and an active or standby (or HA-standby) status for supervisor modules. If the status is either ok or active, you can continue with your configuration.

Note A standby supervisor module reflects the HA-standby status if the HA switchover mechanism is enabled (see the "HA Switchover Characteristics" section on page 15-2). If the warm switchover mechanism is enabled, the standby supervisor module reflects the standby status.

The states through which a switching module progresses is discussed in the "Checking the State of a Module" section on page 17-5.

Obtaining Supervisor Module Statistics

You can view statistics for the supervisor module, such as CPU utilization and NVRAM size, using Fabric Manager. To view supervisor module statistics using Fabric Manager, follow these steps:

Step 1 Do one of the following in the Logical Domains pane:

- Expand **SAN** to display a list of all switches in the SAN.
- Click one of the fabrics to display a list of switches for that fabric.
- Click a VSAN to display a list of switches for that VSAN.

Step 2 Expand **Switches** and select **Supervisor Statistics** in the Physical Attributes pane.

You see the supervisor statistics for each switch in the Information pane.

Checking the State of a Module

The switching module goes through a testing and an initializing stage before displaying an `ok` status. Table 17-3 describes the possible states in which a module can exist.

Table 17-3 Module States

Module Status Output	Description
`powered up`	The hardware has electrical power. When the hardware is powered up, the software begins booting.
`testing`	The switching module has established connection with the supervisor module and the switching module is performing bootup diagnostics.
`initializing`	The diagnostics have completed successfully and the configuration is being downloaded.
`failure`	The switch detects a switching module failure upon initialization and automatically attempts to power-cycle the module three times. After the third attempt it continues to display a failed state.
`ok`	The switch is ready to be configured.
`power-denied`	The switch detects insufficient power for a switching module to power up.
`active`	This module is the active supervisor module and the switch is ready to be configured.
`HA-standby`	The HA switchover mechanism is enabled on the standby supervisor module (see the "HA Switchover Characteristics" section on page 15-2).
`standby`	The warm switchover mechanism is enabled on the standby supervisor module (see the "HA Switchover Characteristics" section on page 15-2).

To view the state of a module from Device Manager, choose **Physical > Modules**. The dialog box displays the status of every module.

Reloading Modules

You can reload the entire switch, reset specific modules in the switch, or reload the image on specific modules in the switch.

Reloading a Switch

To reload a switch using Fabric Manager, follow these steps:

Step 1 Do one of the following in the Logical Domains pane:

- Click **SAN** to display a list of all switches in the SAN.
- Click one of the fabrics to display a list of switches for that fabric.
- Click a VSAN to display a list of switches for that VSAN.

Step 2 Expand **Switches** and select **Hardware** in the Physical Attributes pane (see Figure 17-2).

You see a list of modules contained in the selected switches.

Step 3 Click the **Card Module Status** tab.

You see the information shown in Figure 17-2.

Figure 17-2 Card Module Staus Tab

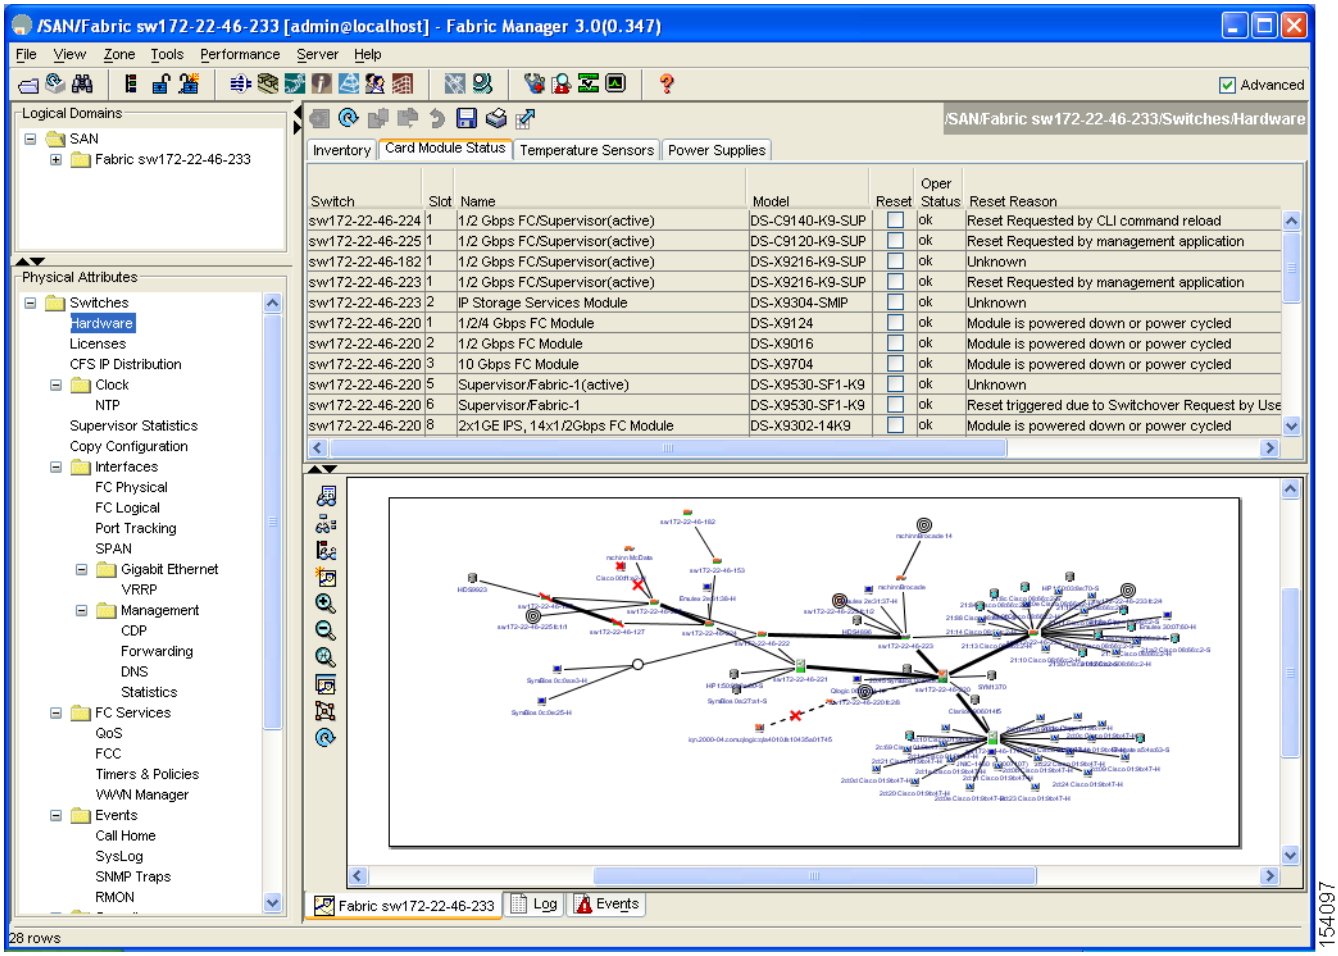

Step 4 Check the **Reset** check box in the row of the switch to reload.

Step 5 Click the **Apply Changes** icon.

Power Cycling Modules

To power cycle any module using Fabric Manager, follow these steps:

Step 1 Do one of the following in the Logical Domains pane:

- Click **SAN** to display a list of all switches in the SAN.
- Click one of the fabrics to display a list of switches for that fabric.
- Click a VSAN to display a list of switches for that VSAN.

Step 2 Expand **Switches** and select **Hardware** from the Physical Attributes pane.

Step 3 Click the **Card Module Status** tab.

Step 4 Check the **Reset** check box in the row for the module(s) you want to reset.

Step 5 Click the **Apply Changes** icon.

Preserving Module Configuration

Use the "copy running-config to startup-config" procedure to save the new configuration into nonvolatile storage. Once this procedure is complete, the running and the startup copies of the configuration are identical.

To preserve the module configuration using Fabric Manager, follow these steps:

Step 1 Do one of the following in the Logical Domains pane:

- Click **SAN** to display a list of all switches in the SAN.
- Click one of the fabrics to display a list of switches for that fabric.
- Click a VSAN to display a list of switches for that VSAN.

Step 2 Expand **Switches** and select **Copy Configuration** in the Physical Attributes pane.

You see a list of switches (see Figure 17-3).

Figure 17-3 List of Switches Available to Copy

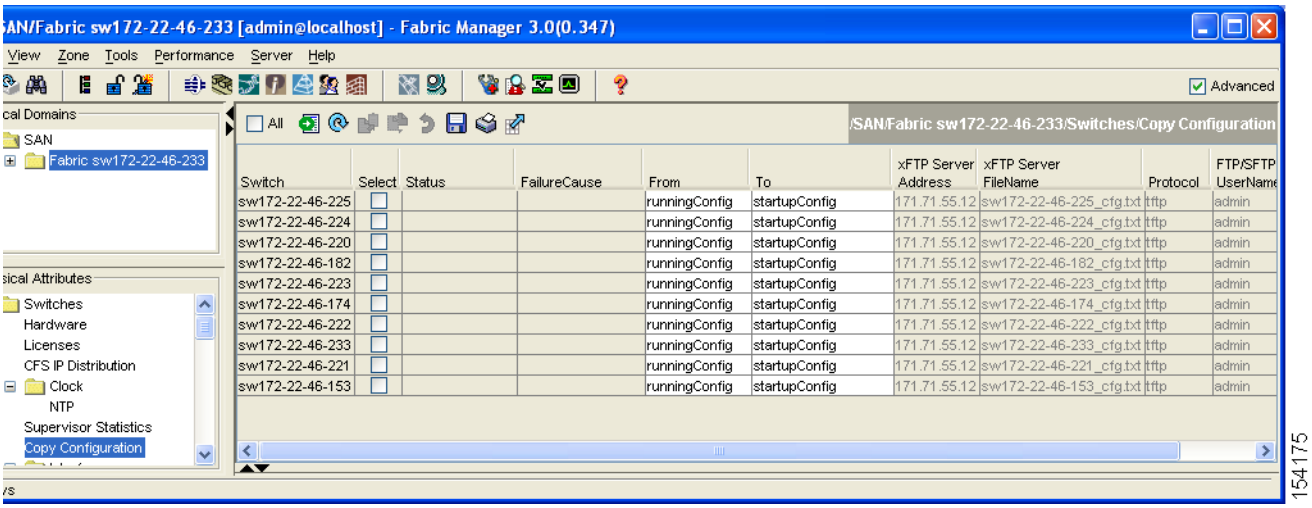

Step 3 Click individual **Select** check boxes for switch configurations to copy.

Step 4 In the From column, ensure that **runningConfig** is selected.

Step 5 In the To column, ensure that **startupConfig** is selected.

Step 6 Click the **Apply Changes** icon.

Table 17-4 displays various scenarios when module configurations are preserved or lost.

Table 17-4 Switching Module Configuration Status

Scenario	Consequence
A particular switching module is removed and the copy running-config startup-config procedure is performed again.	The configured module information is lost.
A particular switching module is removed and the same switching module is replaced before the copy running-config startup-config procedure is performed again.	The configured module information is preserved.
A particular switching module is removed and replaced with the same type switching module, and a module reload procedure is performed on that module.	The configured module information is preserved.
A particular switching module is reloaded when a module reload procedure is performed.	The configured module information is preserved.
A particular switching module is removed and replaced with a different type of switching module. For example, a 16-port switching module is replaced with a 32-port switching module. Sample scenario: 1. The switch currently has a 16-port switching module and the startup and running configuration files are the same. 2. You replace the 16-port switching module in the switch with a 32-port switching module. 3. Next, you remove the 32-port switching module and replace it with the same 16-port switching module referred to in Step 1. 4. You reload the switch.	The configured module information is lost from the running configuration. The default configuration is applied. The configured module information remains in startup configuration until a copy running-config startup-config procedure is performed again. Sample response: 1. The switch uses the 16-port switching module and the present configuration is saved in nonvolatile storage. 2. The factory default configuration is applied. 3. The factory default configuration is applied. 4. The configuration saved in nonvolatile storage referred to in Step 1 is applied.

Powering Off Switching Modules

By default, all switching modules are in the power up state.

To power off a module using Fabric Manager, follow these steps:

Step 1 Do one of the following in the Logical Domains pane:

- Click **SAN** to display a list of all switches in the SAN.
- Click one of the fabrics to display a list of switches for that fabric.
- Click a VSAN to display a list of switches for that VSAN.

Step 2 Expand **Switches** and select **Hardware** in the Physical Attributes pane.

You see a list of modules contained in the selected switches.

Step 3 Select **off** from the drop-down list in the row for the module(s) you want to power off.

Step 4 Click the **Apply Changes** icon.

Note To power on a module, repeat Steps 1-4 but select **on** in Step 3.

Identifying Module LEDs

Table 17-5 describes the LEDs for the Cisco MDS 9200 Series integrated supervisor modules.

Table 17-5 LEDs for the Cisco MDS 9200 Series Supervisor Modules

LED	Status	Description
Status	Green	All diagnostics pass. The module is operational (normal initialization sequence).
	Orange	One of the following applies: • The module is booting or running diagnostics (normal initialization sequence). • The inlet air temperature of the system has exceeded the maximum system operating temperature limit (a minor environmental warning). To ensure maximum product life, you should immediately correct the environmental temperature and restore the system to normal operation.
	Red	One of the following applies: • The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. • The inlet air temperature of the system has exceeded the safe operating temperature limits of the card (a major environmental warning). The card has been shut down to prevent permanent damage. The system will be shut down after two minutes if this condition is not cleared.

Table 17-5 LEDs for the Cisco MDS 9200 Series Supervisor Modules (continued)

LED	Status	Description
Speed	On	2-Gbps mode and beacon mode disabled.
	Off	1-Gbps mode and beacon mode disabled.
	Flashing	Beacon mode enabled.
Link	Solid green	Link is up.
	Solid yellow	Link is disabled by software.
	Flashing yellow	A fault condition exists.
	Off	No link.

Table 17-6 describes the LEDs for the Cisco MDS 9200 Series interface module.

Table 17-6 LEDs on the Cisco MDS 9200 Series Interface Module

LED	Status	Description
Status	Green	All diagnostics pass. The module is operational (normal initialization sequence).
	Orange	One of the following applies: • The module is booting or running diagnostics (normal initialization sequence). • The inlet air temperature of the system has exceeded the maximum system operating temperature limit (a minor environmental warning). To ensure maximum product life, you should immediately correct the environmental temperature and restore the system to normal operation.
	Red	One of the following applies: • The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. • The inlet air temperature of the system has exceeded the safe operating temperature limits of the card (a major environmental warning). The card has been shut down to prevent permanent damage.
System	Green	All chassis environmental monitors are reporting OK.
	Orange	One of the following applies: • The power supply failed or the power supply fan failed. • Incompatible power supplies are installed. • The redundant clock failed.
	Red	The temperature of the supervisor module exceeded the major threshold.
MGMT 10/100 Ethernet Link LED	Green	Link is up.
	Off	No link.

Table 17-6 LEDs on the Cisco MDS 9200 Series Interface Module (continued)

LED	Status	Description
MGMT 10/100 Ethernet Activity LED	Green	Traffic is flowing through port.
	Off	No link or no traffic.

Table 17-7 describes the LEDs for the 16-port and 32-port switching modules, and the 4-port, 12-port, 24-port, and 48-port Generation 2 switching modules.

Table 17-7 LEDs for the Cisco MDS 9000 Family Fibre Channel Switching Modules

LED	Status	Description
Status	Green	All diagnostics pass. The module is operational (normal initialization sequence).
	Orange	One of the following applies: • The module is booting or running diagnostics (normal initialization sequence). • The inlet air temperature of the system has exceeded the maximum system operating temperature limit (a minor environmental warning). To ensure maximum product life, you should immediately correct the environmental temperature and restore the system to normal operation.
	Red	One of the following applies: • The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. • The inlet air temperature of the system has exceeded the safe operating temperature limits of the card (a major environmental warning). The card has been shut down to prevent permanent damage.
Speed	On	2-Gbps mode.
	Off	1-Gbps mode.
Link	Solid green	Link is up.
	Steady flashing green	Link is up (beacon used to identify port).
	Intermittent flashing green	Link is up (traffic on port).
	Solid yellow	Link is disabled by software.
	Flashing yellow	A fault condition exists.
	Off	No link.

The LEDs on the supervisor module indicate the status of the supervisor module, power supplies, and the fan module. Table 17-8 provides more information about these LEDs.

Table 17-8 LEDs for the Cisco MDS 9500 Series Supervisor Modules

LED	Status	Description
Status	Green	All diagnostics pass. The module is operational (normal initialization sequence).
	Orange	One of the following applies: • The module is booting or running diagnostics (normal initialization sequence). • An over temperature condition has occurred (a minor threshold has been exceeded during environmental monitoring).
	Red	One of the following applies: • The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. • An over temperature condition occurred (a major threshold was exceeded during environmental monitoring).
System[1]	Green	All chassis environmental monitors are reporting OK.
	Orange	One of the following applies: • The power supply has failed or the power supply fan has failed. • Incompatible power supplies are installed. • The redundant clock has failed.
	Red	The temperature of the supervisor module major threshold has been exceeded.
Active	Green	The supervisor module is operational and active.
	Orange	The supervisor module is in standby mode.
Pwr Mgmt[1]	Green	Sufficient power is available for all modules.
	Orange	Sufficient power is not available for all modules.
MGMT 10/100 Ethernet Link LED	Green	Link is up.
	Off	No link.
MGMT 10/100 Ethernet Activity LED	Green	Traffic is flowing through port.
	Off	No link or no traffic.
CompactFlash	Green	The external CompactFlash card is being accessed.
	Off	No activity.

1. The System and Pwr Mgmt LEDs on a redundant supervisor module are synchronized to the active supervisor module.

Default Settings

Table 17-9 lists the default settings for the supervisor module.

Table 17-9 Default Supervisor Module Settings

Parameters	Default
Administrative connection	Serial connection.
Global switch information	• No value for system name. • No value for system contact. • No value for location.
System clock	No value for system clock time.
In-band (VSAN 1) interface	IP address, subnet mask, and broadcast address assigned to the VSAN are set to 0.0.0.0.

Cisco Systems

Send documentation comments to mdsfeedback-doc@cisco.com

PART 3

Switch Configuration

CHAPTER 18

Configuring Interfaces

A switch's main function is to relay frames from one data link to another. To do that, the characteristics of the interfaces through which the frames are received and sent must be defined. The configured interfaces can be Fibre Channel interfaces, Gigabit Ethernet interfaces, the management interface (mgmt0), or VSAN interfaces.

This chapter describes the basic interface configuration to get your switch up and running. It includes the following sections:

Note See Chapter 11, "Initial Configuration," and Chapter 46, "Configuring IP Services," for more information on configuring mgmt0 interfaces.

Note See the "Configuring Gigabit Ethernet Interfaces for IPv4" section on page 47-4 for more information on configuring Gigabit Ethernet interfaces.

Tip Before you begin configuring the switch, ensure that the modules in the chassis are functioning as designed. See the "Verifying the Module Status" section on page 11-2.

Fibre Channel Interfaces

The 32-port switching module guidelines apply to the following hardware:

- The 32-port, 2-Gbps or 1-Gbps switching module
- The Cisco MDS 9140 Switch

When configuring these host-optimized ports, the following port mode guidelines apply:

- You can configure only the first port in each 4-port group (for example, the first port in ports 1-4, the fifth port in ports 5-8 and so on) as an E port. If the first port in the group is configured as an E port, the other three ports in each group (ports 2-4, 6-8 and so on) are not usable and remain shutdown.
- If any of the other three ports are enabled, you cannot configure the first port as an E port. The other three ports continue to remain enabled.
- The auto mode is not allowed in a 32-port switching module or the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).
- The default port mode is Fx (Fx negotiates to F or FL) for 32-port switching modules and the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).
- The 32-port switching module does not support FICON.

Note We recommend that you configure your E ports on a 16-port switching module. If you must configure an E port on a 32-port host optimized switching module, the other three ports in that four port group cannot be used.

Note In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the 9120 switch and 8 ports in the 9140 switch) are full line rate like the 16-port switching module. The other ports (16 ports in the 9120 switch and 32 ports in the 9140 switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.

This section describes Fibre Channel interface characteristics, including (but not limited to) modes, frame encapsulation, and SFPs. It includes the following topics:

About Interface Modes

Each physical Fibre Channel interface in a switch may operate in one of several port modes: E port, F port, FL port, TL port, TE port, SD port, ST port, and B port (see Figure 18-1). Besides these modes, each interface may be configured in auto or Fx port modes. These two modes determine the port type during interface initialization.

Figure 18-1 Cisco MDS 9000 Family Switch Port Modes

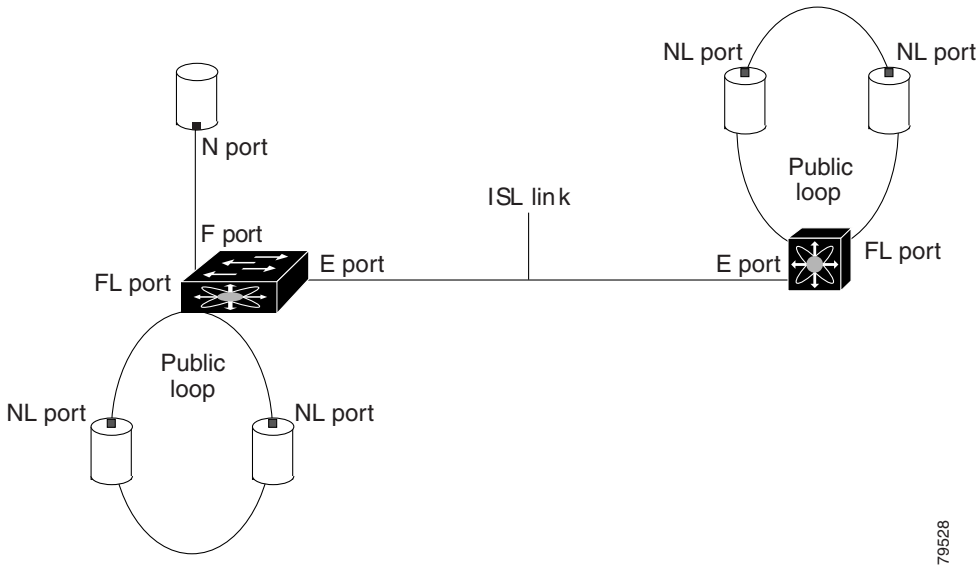

Note Interfaces are created in VSAN 1 by default. See Chapter 23, "Configuring and Managing VSANs."

Each interface has an associated administrative configuration and an operational status:

- The administrative configuration does not change unless you modify it. This configuration has various attributes that you can configure in administrative mode.
- The operational status represents the current status of a specified attribute like the interface speed. This status cannot be changed and is read-only. Some values may not be valid when the interface is down (for example, the operational speed).

A brief description of each interface mode follows.

E Port

In expansion port (E port) mode, an interface functions as a fabric expansion port. This port may be connected to another E port to create an Inter-Switch Link (ISL) between two switches. E ports carry frames between switches for configuration and fabric management. They serve as a conduit between switches for frames destined to remote N ports and NL ports. E ports support class 2, class 3, and class F service.

An E port connected to another switch may also be configured to form a PortChannel (see Chapter 21, "Configuring PortChannels").

F Port

In fabric port (F port) mode, an interface functions as a fabric port. This port may be connected to a peripheral device (host or disk) operating as an N port. An F port can be attached to only one N port. F ports support class 2 and class 3 service.

FL Port

In fabric loop port (FL port) mode, an interface functions as a fabric loop port. This port may be connected to one or more NL ports (including FL ports in other switches) to form a public arbitrated loop. If more than one FL port is detected on the arbitrated loop during initialization, only one FL port becomes operational and the other FL ports enter nonparticipating mode. FL ports support class 2 and class 3 service.

Note FL port mode is not supported on 4-port 10-Gbps switching module interfaces.

TL Port

In translative loop port (TL port) mode, an interface functions as a translative loop port. It may be connected to one or more private loop devices (NL ports). TL ports are specific to Cisco MDS 9000 Family switches and have similar properties as FL ports. TL ports enable communication between a private loop device and one of the following devices:

- A device attached to any switch on the fabric
- A device on a public loop anywhere in the fabric
- A device on a different private loop anywhere in the fabric
- A device on the same private loop

TL ports support class 2 and class 3 services.

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop (see the "Default Settings" section on page 18-18 and the "About TL Port ALPA Caches" section on page 18-11).

Tip We recommend configuring devices attached to TL ports in zones that have up to 64 zone members.

Note TL port mode is not supported on Generation 2 switching module interfaces.

TE Port

In trunking E port (TE port) mode, an interface functions as a trunking expansion port. It may be connected to another TE port to create an Enhanced ISL (EISL) between two switches. TE ports are specific to Cisco MDS 9000 Family switches. They expand the functionality of E ports to support the following:

- VSAN trunking
- Transport quality of service (QoS) parameters
- Fibre Channel trace (fctrace) feature

In TE port mode, all frames are transmitted in EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link. This feature is referred to as trunking in the Cisco MDS 9000 Family (see Chapter 20, "Configuring Trunking"). TE ports support class 2, class 3, and class F service.

SD Port

In SPAN destination port (SD port) mode, an interface functions as a switched port analyzer (SPAN). The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic that passes though a Fibre Channel interface. This monitoring is done using a standard Fibre Channel analyzer (or a similar switch probe) that is attached to an SD port. SD ports do not receive frames, they merely transmit a copy of the source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports (see Chapter 56, "Monitoring Network Traffic Using SPAN").

ST Port

In the SPAN tunnel port (ST port) mode, an interface functions as an entry point port in the source switch for the RSPAN Fibre Channel tunnel. The ST port mode and the remote SPAN (RSPAN) feature are specific to switches in the Cisco MDS 9000 Family. When configured in ST port mode, the interface cannot be attached to any device, and thus cannot be used for normal Fibre Channel traffic (see the "Default Settings" section on page 56-13).

Fx Port

Interfaces configured as Fx ports can operate in either F port or FL port mode. The Fx port mode is determined during interface initialization depending on the attached N port or NL port. This administrative configuration disallows interfaces to operate in any other mode—for example, preventing an interface to connect to another switch.

B Port

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as the Cisco PA-FC-1G Fibre Channel port adapter, implement a bridge port (B port) model to connect geographically dispersed fabrics. This model uses B ports as described in the T11 Standard FC-BB-2.

Figure 47-1 on page 47-2 depicts a typical SAN extension over an IP network.

If an FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled (see Chapter 47, "Configuring IP Storage").

Auto Mode

Interfaces configured in auto mode can operate in one of the following modes: F port, FL port, E port, or TE port. The port mode is determined during interface initialization. For example, if the interface is connected to a node (host or disk), it operates in F port or FL port mode depending on the N port or NL port mode. If the interface is attached to a third-party switch, it operates in E port mode. If the interface is attached to another switch in the Cisco MDS 9000 Family, it may become operational in TE port mode (see Chapter 20, "Configuring Trunking").

TL ports and SD ports are not determined during initialization and are administratively configured.

Note Fibre Channel interfaces on Storage Services Modules (SSMs) cannot be configured in auto mode.

Configuring Interface Modes

To configure the interface mode using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **General** tab.

Step 3 Click **Mode Admin**. Set the desired interface mode from the Admin drop-down menu.

Step 4 Optionally, set other configuration parameters using the other tabs.

Step 5 Click **Apply Changes** icon.

Configuring Port Speeds

By default, the port speed for an interface is automatically calculated by the switch.

Caution Changing the port speed is a disruptive operation.

To configure the interface mode using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **General** tab.

Step 3 Click **Speed Admin**. Set the desired speed from the drop-down menu.

The number indicates the speed in megabits per second (Mbps). You can set the speed to 1-Gbps, 2-Gbps, 4-Gbps, or **auto** (default).

Step 4 Click **Apply Changes** icon.

About Frame Encapsulation

You can set the frame format to EISL for all frames transmitted by an interface in SD port mode. If you sent the frame encapsulation to EISL, all outgoing frames are transmitted in the EISL frame format, irrespective of the SPAN source(s). See the "Monitoring Network Traffic Using SPAN" section on page 56-1.

Refer to the *Cisco MDS 9000 Family CLI Configuration Guide* to configure frame encapsulation on an interface.

About Receive Data Field Size

You can also configure the receive data field size for Fibre Channel interfaces. If the default data field size is 2112 bytes, the frame length will be 2148 bytes.

Configuring Receive Data Field Size

To configure data field size using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** and then select **FC Physical**. You see the interface configuration in the Information pane.

Step 2 Click the **Other** tab and set the RxDataFieldSize field shown in blue in Figure 18-2.

Figure 18-2 Changing Rx Data Size

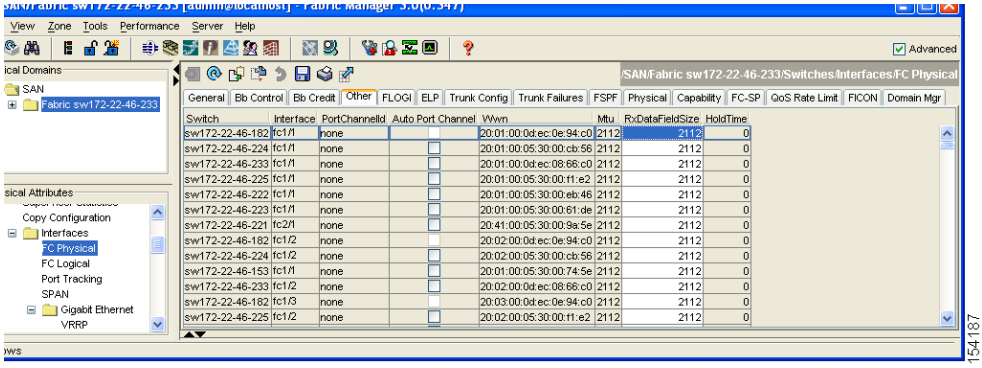

Step 3 Optionally, set other configuration parameters using the other tabs.

Step 4 Click **Apply Changes**.

Configuring the Beacon Mode

By default, the beacon mode is disabled on all switches. The beacon mode is indicated by a flashing green light that helps you identify the physical location of the specified interface.

Configuring the beacon mode has no effect on the operation of the interface.

To enable beacon mode for a specified interface or range of interfaces using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** and then select **Gigabit Ethernet**. You see the interface configuration in the Information pane.

Step 2 Enable the Beacon Mode option for the selected switch.

Step 3 Click **Apply Changes**.

Note The flashing green light turns on automatically when an external loopback is detected that causes the interfaces to be isolated. The flashing green light overrides the beacon mode configuration. The state of the LED is restored to reflect the beacon mode configuration after the external loopback is removed.

About Bit Error Thresholds

The bit error rate threshold is used by the switch to detect an increased error rate before performance degradation seriously affects traffic.

The bit errors can occur for the following reasons:

- Faulty or bad cable.
- Faulty or bad GBIC or SFP.
- GBIC or SFP is specified to operate in 1Gbps but is used for 2Gbps.
- GBIC or SFP is specified to operate in 2 Gbps but is used for 4Gbps.
- Short haul cable used for long haul or long haul cable is used for short haul.
- Momentary sync loss.
- Loose cable connection at one or both ends.
- Improper GBIC or SFP connection at one or both ends.

A bit error rate threshold is detected when 15 error bursts occur in a 5-minute period. By default, the switch disables the interface when the threshold reached. You can reenable the interface.

You can configure the switch to not disable an interface when the threshold is crossed. By default, the threshold disables the interface.

Refer to the *Cisco MDS 9000 Family CLI Configuration Guide* to disable the bit error threshold for an interface.

Note Regardless of disabling the bit-error threshold for an interface, the switch generates a syslog message when bit error threshold events are detected.

About SFP Transmitter Types

The small form-factor pluggable (SFP) hardware transmitters are identified by their acronyms when displayed. Table 18-1 defines the acronyms used for SFPs (see the "Displaying SFP Transmitter Types" section on page 18-9).

Table 18-1 SFP Transmitter Acronym Definitions

Definition	Acronym
Standard transmitters defined in the GBIC specifications	
short wavelength laser	swl
long wavelength laser	lwl
long wavelength laser cost reduced	lwcr
electrical	elec
Extended transmitters assigned to Cisco-supported SFPs	
CWDM-1470	c1470
CWDM-1490	c1490
CWDM-1510	c1510
CWDM-1530	c1530
CWDM-1550	c1550
CWDM-1570	c1570
CWDM-1590	c1590
CWDM-1610	c1610

Displaying SFP Transmitter Types

To show the SFP types for an interface using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** and then select **FC Physical**. You see the interface configuration in the Information pane.

Step 2 Click the **Physical** tab to see the transmitter type for the selected interface.

TL Ports for Private Loops

Private loops require setting the interface mode to TL. This section describes TL ports and includes the following sections:

- About TL Ports, page 18-10
- Configuring TL Ports, page 18-11
- About TL Port ALPA Caches, page 18-11

About TL Ports

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop.

The legacy devices are used in Fibre Channel networks and devices outside the loop may need to communicate with them.The communication functionality is provided through TL ports. See the "About Interface Modes" section on page 18-3.

Follow these guidelines when configuring private loops:

- A maximum of 64 fabric devices can be proxied to a private loop.
- Fabric devices must be in the same zone as private loop devices to be proxied to the private loop.
- Each private device on a TL port may be included in a different zone.
- All devices on the loop are treated as private loops. You cannot mix private and public devices on the loop if the configured port mode is TL.
- The only FC4-type supported by TL ports is SCSI (FCP).
- Communication between a private initiator to a private target on the same private loop does not invoke TL port services.

Table 18-2 lists the TL port translations supported in Cisco MDS 9000 Family switches. Figure 18-3 shows examples of TL port translation support.

Table 18-2 Supported TL Port Translations

Translation from	Translation to	Example
Private initiator	Private target	From I1 to T1 or vice versa
Private initiator	Public target — N port	From I1 to T2 or vice versa
Private initiator	Public target — NL port	From I4 to T3 or vice versa
Public initiator — N port	Private target	From I2 to T1 or vice versa
Public initiator — NL port	Private target	From I3 to T1 or vice versa

Figure 18-3 TL Port Translation Support Examples

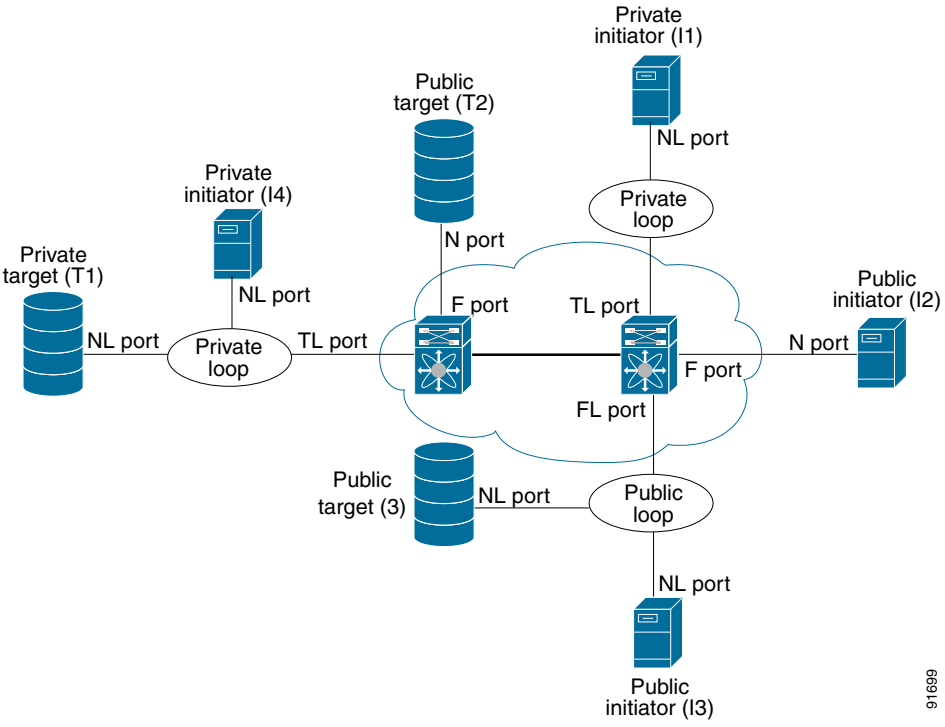

Configuring TL Ports

To configure the TL interface mode using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** and then select **FC Physical**. You see the interface configuration in the Information pane.

Step 2 Choose the **General** tab and click **Mode Admin.**

Step 3 Set the Mode Admin drop-down menu to **TL**.

Step 4 Optionally, set other configuration parameters using the other tabs.

Step 5 Click **Apply Changes**.

About TL Port ALPA Caches

Although TL ports cannot be automatically configured, you can manually configure entries in arbitrated loop physical address (ALPA) caches. Generally, ALPA cache entries are automatically populated when an ALPA is assigned to a device. Each device is identified by its port world wide name (pWWN). When a device is allocated an ALPA, an entry for that device is automatically created in the ALPA cache.

A cache contains entries for recently allocated ALPA values. These caches are maintained on various TL ports. If a device already has an ALPA, the Cisco SAN-OS software attempts to allocate the same ALPA to the device each time. The ALPA cache is maintained in persistent storage and saves information across switch reboots. The maximum cache size is 1000 entries. If the cache is full, and a new ALPA is allocated, the Cisco SAN-OS software discards an inactive cache entry (if available) to make space for the new entry.

See the "Default Settings" section on page 18-18 for more information on TL ports. Refer to the *Cisco MDS 9000 Family CLI Configuration Guide* to manage the TL Port ALPA cache.

Buffer Credits

Fibre Channel interfaces use buffer credits to ensure all packets are delivered to their destination. This section describes the different buffer credits available on the Cisco MDS Family switches and includes the following topics:

About Buffer-to-Buffer Credits

Buffer-to-buffer credits (BB_credits) are a flow control mechanism to ensure that FC switches do not run out of buffers, because switches must not drop frames. BB_credits are negotiated on a per-hop basis.

The receive BB_credit (`fcrxbbcredit`) value may be configured for each FC interface. In most cases, you do not need to modify the default configuration.

Note The receive BB_credit values depend on the module type and the port mode.
For 16-port switching modules and full rate ports, the default value is 16 for Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required.
For 32-port switching modules and host-optimized ports, the default value is 12 for Fx, E, and TE modes. These values cannot be changed.

Note In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the 9120 switch and 8 ports in the 9140 switch) are full line rate like the 16-port switching module. The other ports (16 ports in the 9120 switch and 32 ports in the 9140 switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.

Configuring Buffer-to-Buffer Credits

To configure BB_credits for a Fibre Channel interface using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** and then select **FC Physical**. You see the interface configuration in the Information pane.

Step 2 Choose the **Bb Credit** tab.

You see the buffer credits.

Step 3 Set any of the buffer-to-buffer credits for an interface.

Step 4 Click **Apply Changes**.

About Performance Buffers

Regardless of the configured receive BB_credit value, additional buffers, called performance buffers, improve switch port performance. Instead of relying on the built-in switch algorithm, you can manually configure the performance buffer value for specific applications (for example, forwarding frames over FCIP interfaces).

For each physical Fibre Channel interface in any switch in the Cisco MDS 9000 Family, you can specify the amount of performance buffers allocated in addition to the configured receive BB_credit value.

The default performance buffer value is 0. If you set the performance buffer value to 0, the built-in algorithm is used. If you do not specify the performance buffer value, 0 is automatically used.

Configuring Performance Buffers

To configure performance buffers for a Fibre Channel interface using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** and then select **FC Physical**. You see the interface configuration in the Information pane.

Step 2 Click the **BB Credit** tab

You see performance buffer information in the columns Perf Bufs Admin and Perf Bufs Oper.

Step 3 Set the performance buffers for an interface.

Step 4 Click **Apply Changes**.

About Extended BB_credits

The BB_credits feature allows you to configure up to 255 receive buffers. This number is insufficient for long haul links. To facilitate BB_credits for long haul links, you can use the extended BB_credits flow control mechanism. This feature allows you to configure up to 3,500 receive BB_credits on a Fibre Channel port.

This section includes the following topics:

Extended BB_credits on Generation 1 Switching Modules

The BB_credits feature allows you to configure up to 255 receive buffers on Generation 1 switching modules. To facilitate BB_credits for long haul links, you can configure up to 3,500 receive BB_credits on a Fibre Channel port on a Generation 1 switching module.

To use this feature on Generation 1 switching modules, you must meet the following requirements:

- Obtain the ENTERPRISE_PKG license (see Chapter 10, "Obtaining and Installing Licenses").
- Configure this feature in any port of the full-rate 4-port group in either the Cisco MDS 9216i Switch or in the MPS-14/2 module (see Figure 18-4).

Figure 18-4 Port Group Support for the Extended BB_Credits Feature

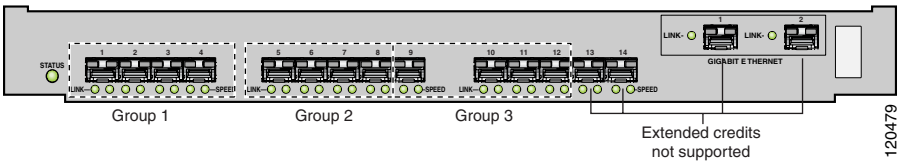

The port groups that support extended credit configurations are as follows.

- Any one port in ports 1 to 4 (identified as Group 1 in Figure 18-1).
- Any one port in ports 5 to 8 (identified as Group 2 in Figure 18-1).
- Any one port in ports 9 to 12 (identified as Group 3 in Figure 18-1).

Note The last two Fibre Channel ports (port 13 and port 14) and the two Gigabit Ethernet ports do not support the extended BB_credits feature (see Figure 18-1).

- Explicitly enable this feature in the required Cisco MDS switch.
- Disable the remaining three ports in the 4-port group if you need to assign more than 2,400 BB_credits to the first port in the port group.
  - If you assign less than 2,400 extended BB_credits to any one port in a port group, the remaining three ports in that port group can retain up to 255 BB_credits based on the port mode.

    Note The receive BB_credit value for the remaining three ports depends on the port mode. The default value is 16 for the Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required without exceeding the maximum value of 255 BB_credits.

  - If you assign more than 2,400 (up to a maximum of 3,500) extended BB_credits to the port in a port group, you must disable the other three ports.
- Be aware that changing the BB_credit value results in the port being disabled and then reenabled.
- Disable (explicitly) this feature if you need to nondisruptively downgrade to Cisco SAN-OS Release 1.3 or earlier. When you disable this feature, the existing extended BB_credit configuration is completely erased.

Note The extended BB_credit configuration takes precedence over the receive BB_credit and performance buffer configurations.

Extended BB_credits on Generation 2 Switching Modules

To use this feature on Generation 2 switching modules, you must meet the following requirements:

- Obtain the Enterprise package (ENTERPRISE_PKG) license (see Chapter 10, "Obtaining and Installing Licenses").
- Configure this feature in any port on a Generation 2 switch module. See the Extended BB_Credits, page 19-10 for more information on extended BB_credits on Generation 2 switching modules.

Configuring Extended BB_credits

To configure extended BB_credits for an MDS-14/2 interface, for a Generation 2 switching module interface, or for an interface in a Cisco MDS 9216i switch using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** and then select **FC Physical**. You see the interface configuration in the Information pane.

Step 2 Choose the **BB Credit** tab.

Step 3 In the Extended column, set the extended BB_credits for the selected interface.

Step 4 Click **Apply Changes**.

Management Interfaces

You can remotely configure the switch through the management interface (mgmt0). To configure a connection remotely, you must configure the IP parameters (IP address, subnet mask, and default gateway) so that the switch is reachable.

This section describes the management interfaces and includes the following topics:

- About Management Interfaces, page 18-15
- Configuring Management Interfaces, page 18-16

About Management Interfaces

Before you begin to configure the management interface manually, obtain the switch's IP address and IP subnet mask.

The management port (mgmt0) is autosensing and operates in full duplex mode at a speed of 10/100 Mbps. The speed and mode cannot be configured.

Note You need to explicitly configure a default gateway to connect to the switch and send IP packets or add a route for each subnet.

Configuring Management Interfaces

To configure the mgmt0 Ethernet interface using Fabric Manager, follow these steps:

Step 1 Select a VSAN in the Logical Domains pane.

Step 2 Expand **Switches > Interfaces** and then select **Management**.

You see the interface configuration in the Information pane.

Step 3 Click the **General** tab.

Step 4 Set the IP Address/Mask field.

Step 5 Set Admin to **up**.

Step 6 Optionally, set other configuration parameters using the other tabs.

Step 7 Click **Apply Changes**.

VSAN Interfaces

VSANs apply to Fibre Channel fabrics and enable you to configure multiple isolated SAN topologies within the same physical infrastructure. You can create an IP interface on top of a VSAN and then use this interface to send frames to this VSAN. To use this feature, you must configure the IP address for this VSAN. VSAN interfaces cannot be created for nonexisting VSANs.

This section describes VSAN interfaces and includes the following topics:

About VSAN Interfaces

Follow these guidelines when creating or deleting VSAN interfaces:

- Create a VSAN before creating the interface for that VSAN. If a VSAN does not exist, the interface cannot be created.
- Create the interface VSAN—it is not created automatically.
- If you delete the VSAN, the attached interface is automatically deleted.
- Configure each interface only in one VSAN.

Tip After configuring the VSAN interface, you can configure an IP address or Virtual Router Redundancy Protocol (VRRP) feature (see Chapter 46, “Configuring IP Services”).

Configuring VSAN Interfaces

To create a VSAN interface using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** and then select **Management**.

You see the interface configuration in the Information pane (see Figure 18-5).

Figure 18-5 General Management Tab

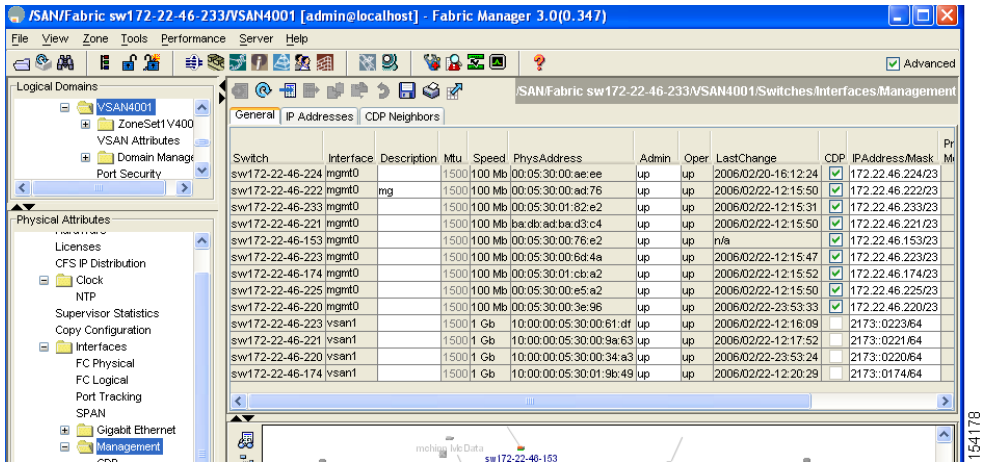

Step 2 Click **Create Row**.

You see the Create Interface dialog box (see Figure 18-6).

Figure 18-6 Create Interface

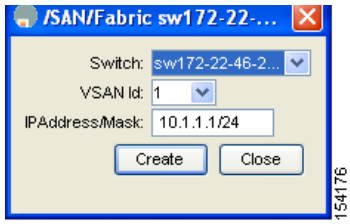

Step 3 Select the switch and VSAN ID for which you want to configure a VSAN interface.

Note You can only create a VSAN interface for an existing VSAN. If the VSAN does not exist, you cannot create a VSAN interface for it.

Step 4 Set IPAddress/Mask to the IP address and subnet mask for the new VSAN interface.

Step 5 Click **Create** to create the VSAN interface or click **Close** to close the dialog box without creating the VSAN interface.

Default Settings

Table 18-3 lists the default settings for interface parameters.

Table 18-3 Default Interface Parameters

Parameters	Default
Interface mode	Auto.
Interface speed	Auto.
Administrative state	Shutdown (unless changed during initial setup).
Trunk mode	On (unless changed during initial setup).
Trunk-allowed VSANs	1 to 4093.
Interface VSAN	Default VSAN (1).
Beacon mode	Off (disabled).
EISL encapsulation	Disabled.
Data field size	2112 bytes.

CHAPTER 19

Configuring Generation 2 Switching Modules

The Cisco MDS 9500 Series switches and Cisco MDS 9216A and Cisco MDS 9216i switches support a set of modules called Generation 2 modules. This chapter describes how to configure these modules.

This chapter includes the following sections:

About Generation 2 Modules

The Cisco MDS 9500 Series switches and Cisco MDS 9216A and Cisco MDS 9216i switches support the following Generation 2 modules:

- 48-port 4-Gbps Fibre Channel switching module (part number DS-X9148)
- 24-port 4-Gbps Fibre Channel switching module (part number DS-X9124)
- 12-port 4-Gbps Fibre Channel switching module (part number DS-X9112)
- 4-port 10-Gbps Fibre Channel switching module (part number DS-X9704)
- Supervisor-2 module (Cisco MDS 9500 Series switches only) (part number DS-X9530-SF2-K9)

Note Generation 2 Fibre Channel switching modules are not supported on the Cisco MDS 9216 switch.

For detailed information about the installation and specifications for these modules, refer to the hardware installation guide for your switch.

This section includes the following topics:

Port Groups

Each module has four groups of one or more ports that have a combined bandwidth of up to 12.8 Gbps. Table 19-1 shows the port groups for the Generation 2 Fibre Channel switching modules.

Table 19-1 Bandwidth and Port Groups for Generation 2 Modules

Switching Module	Number of Ports Per Port Group	Bandwidth Per Port Group	Maximum Bandwidth Per Port
48-port 4-Gbps	12	12.8	4-Gbps[2]
24-port 4-Gbps	6	12.8	4-Gbps[1]
12-port 4-Gbps	3	12.8	4-Gbps[2]
4-port 10-Gbps	1	10	10-Gbps[1]

1. Dedicated bandwidth or oversubscribed using shared buffer resources.
2. Dedicated bandwidth with no oversubscription.

Note Port groups are defined by the hardware and consist of sequential ports. For example, ports 1 through 12, ports 13 through 24, ports 25 through 36, and ports 37 through 48 are the port groups on the 48-port 4-Gbps switching modules.

Port Rate Modes

The ports on the 24-port and 48-port 4-Gbps switching modules can support both dedicated and oversubscribed (or shared) bandwidth. By default, all the ports on these modules support 4 Gbps of shared bandwidth. You can dedicate 1 Gbps, 2 Gbps or 4 Gbps on any of the ports in a port group, if your deployment requires it, up to a maximum of 12 Gbps per port group. If enough undedicated bandwidth is available, it is shared over the remaining ports in the port group.

Note All ports on the 12-port 4-Gbps switching module and 4-port 10-Gbps switching module operate only in dedicated mode.

Dedicated Mode

When port rate mode is configured as dedicated, a port is allocated required fabric bandwidth and related resources to sustain line rate traffic at the maximum operating speed configured for the port. In this mode, ports do not use local buffering and all receive buffers are allocated from a global buffer pool (see the "Buffer Groups" section on page 19-5).

All ports in a 24-port 4-Gbps switching module can operate in dedicated mode with a 2-Gbps operating speed. However, if you configure one or more ports to operate in 4-Gbps dedicated mode, some of the other ports in the module would have to operate in shared mode.

All ports in a 48-port 4-Gbps switching module can operate in dedicated mode with a 1-Gbps operating speed. However, if you configure one or more ports to operate in 2-Gbps or 4-Gbps dedicated mode, some of the other ports in the module would have to operate in shared mode.

Table 19-2 show the amount of bandwidth reserved for a configured port speed on 4-Gbps switching modules.

Table 19-2 Bandwidth Reserved for the Port Speeds on 4-Gbps Switching Modules

Configured Speed	Reserved Bandwidth
Auto	4 Gbps
4-Gbps	
Auto with 2-Gbps maximum	2 Gbps
2-Gbps	
1-Gbps	1 Gbps

Note The 4-port 10-Gbps switching module only supports auto speed mode at 10 Gbps.

Shared Mode

When port rate mode is configured as shared, multiple ports share data paths to the switch fabric so that fabric bandwidth and related resources are shared. Often, the available bandwidth to the switch fabric may be less than the negotiated operating speed of a port. Ports in this mode use local buffering for the BB_credit buffers.

All ports in 24-port and 48-port 4-Gbps switching modules operate by default in shared mode and in administrative operating speed auto, which supports 1-Gbps, 2-Gbps, or 4-Gbps traffic. However, it is possible to configure one or more ports in a port group to operate in dedicated mode with either 2-Gbps or 4-Gbps operating speed.

Dynamic Bandwidth Management

On 24-port and 48-port 4-Gbps switching modules, the bandwidth available to each port within a port group can be configured based on the port rate mode and speed configurations. Within a port group, some ports can be configured in dedicated mode while others operate in shared mode. Ports configured in dedicated mode are allocated the required bandwidth to sustain a line rate of traffic at the maximum configured operating speed, and ports configured in shared mode share the available remaining bandwidth within the port group. The bandwidth allocation among the shared mode ports are based on the operational speed of the ports. For example, if four ports operating at speeds 1 Gbps, 1 Gbps, 2 Gbps, and 4 Gbps share bandwidth of 8 Gbps, the ratio of allocation would be 1:1:2:4.

Autosensing

Autosensing speed is enabled on all 4-Gbps switching module interfaces by default. This configuration allows the interfaces operates at speeds of 1 Gbps, 2 Gbps, or 4 Gbps on the 4-Gbps switching modules. When autosensing is enabled for an interface operating in dedicated rate mode, 4-Gbps bandwidth is reserved, even if the port negotiates an operating speed of 1-Gbps or 2-Gbps. To avoid wasting unused bandwidth on 24-port and 48-port 4-Gbps switching modules, you can specify that only 2 Gbps of required bandwidth be reserved, not the default of 4 Gbps. This feature shares the unused bandwidth within the port group provided it does not exceed the rate limit configuration for the port. You can also use this feature for shared rate ports configured for autosensing.

Tip When migrating from Generation 1 switching modules to 4-Gbps switching modules, use autosensing with a maximum bandwidth of 2-Gbps maximum.

Note If you configure an interface for autosensing speed with a maximum bandwidth of 2 Gbps and want to change to the default of 4 Gbps, be sure that there are enough shared resources available to support the configuration on the module.

Note The 4-port 10-Gbps switching module only supports 10-Gbps speed traffic.

Oversubscription

The 24-port and 48-port 4-Gbps switching modules support oversubscription when shared allocation is configured. Table 19-3 describes the bandwidth allocation for oversubscribed interfaces configured in shared mode.

Table 19-3 Bandwidth Allocation for Oversubscribed Interfaces

Switching Module Type	Configured Speed	Reserved Bandwidth (Gbps)	Maximum Bandwidth (Gbps)
24-port	Auto 4 Gbps	1	4
	Auto max 2 Gbps 2 Gbps	0.5	2
	1 Gbps	0.25	1
48-port	Auto 4 Gbps	0.8	4
	Auto max 2 Gbps 2 Gbps	0.4	2
	1 Gbps	0.2	1

Out-of-Service Interfaces

On the 24-port and 48-port 4-Gbps switching module, you might need to allocate all the shared resources for one or more interfaces to another interface in the port group or module. You can take interfaces out of service to release shared resources that are needed for dedicated bandwidth. When an interface is taken out of service, all shared resources are released and made available to the other interface in the port group or module. These shared resources include speed mode, rate mode, BB_credits, and extended BB_credits. All shared resource configurations are returned to their default values when the interface is brought back into service.

If you need to bring an interface back into service, you might disrupt traffic if you need to release shared resources from other interfaces.

Caution If you need to bring an interface back into service, you might disrupt traffic if you need to release shared resources from other interfaces.

Buffer Groups

In the architecture of Generation 2 modules, receive buffers shared by a set of ports are called *buffer groups*. The receive buffer groups are organized in global and local buffer pools.

The receive buffers allocated from the global buffer pool to be shared by a port group are called a *global buffer group*. Global receive buffer pools include the following buffer groups:

- Reserved internal buffers
- Allocated BB_credit buffers for each Fibre Channel interface (user configured or assigned by default)
- Common unallocated buffer pool for BB_credits, if any, to be used for additional BB_credits as needed
- Performance buffers (only used on 12-port 4-Gbps and 4-port 10-Gbps switching modules)

Receive BB_Credit Buffers

Figure 19-1 shows the default BB_credit buffer allocation model for 24-port 4-Gbps switching modules. The minimum BB_credits required to bring up a port is two buffers.

Figure 19-1 BB_Credit Buffer Allocation in 24-port 4-Gbps Switching Modules

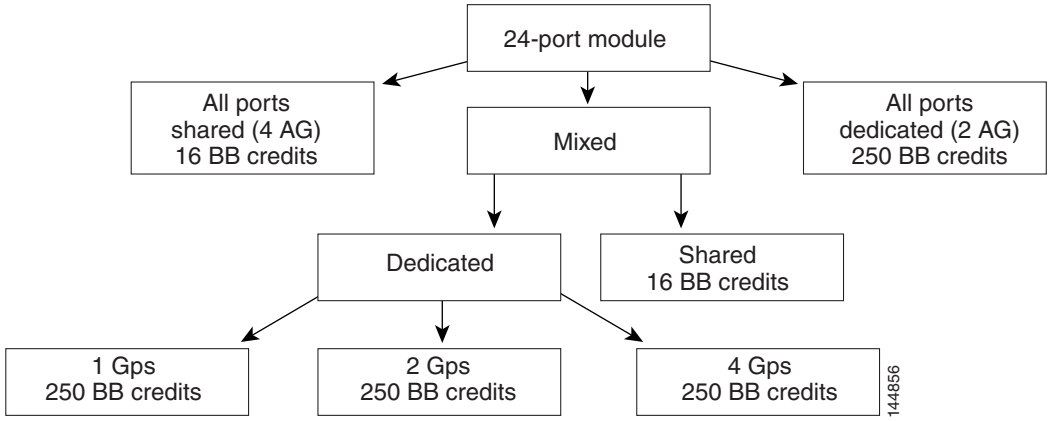

Note The default BB_credit buffer allocation is the same for all port speeds.

48-port 4-Gbps Switching Module BB_Credit Buffers

Table 19-4 lists the BB_credit buffer allocation for 48-port 4-Gbps switching modules.

Table 19-4 48-Port 4-Gbps Switching Module BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Module	BB_Credit Buffers Per Port Defaults			
		Dedicated Rate Mode 4-Gbps Speed		Shared Rate Mode 4-Gbps Speed	
		ISL[1]	Fx Port	ISL[1]	Fx Port
User configurable BB_credit buffers	6000	125	16	16	16

1. ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 24-port 4-Gbps switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- Performance buffers are not supported on this module.

Each port group on the 48-port 4-Gbps switching module consists of 12 ports. The ports in shared rate mode have bandwidth oversubscription of 4:1 by default. However, some configurations of the shared ports in a port group can have maximum bandwidth oversubscription of 5:1 (considering that each port group has 12.8-Gbps bandwidth).

The following example configurations are supported by the 48-port 4-Gbps switching modules:

- Twelve ports with shared rate mode and 4-Gbps speed (4:1 oversubscription) (default)
- One port with dedicated rate mode and 4-Gbps speed plus
  11 ports with shared rate mode and 4-Gbps speed (5:1 oversubscription)
- One port with dedicated rate mode and 4-Gbps speed plus
  11 ports with shared rate mode and 2-Gbps speed (2.5:1 oversubscription)
- Two ports with dedicated rate mode and 2-Gbps speed plus
  10 ports with shared rate mode and 4-Gbps speed (5:1 oversubscription)
- Two ports with dedicated rate mode and 2-Gbps speed plus
  10 ports with shared rate mode and 2-Gbps speed (2.5:1 oversubscription)
- Twelve ports with dedicated rate mode and 1-Gbps speed
- Three ports with dedicated rate mode and 4-Gbps speed plus
  four ports with shared rate mode and 1-Gbps speed plus
  five ports put out-of-service (see Figure 19-2)

Figure 19-2 Example Speed and Rate Configuration on a 48-Port 4-Gbps Switching Module

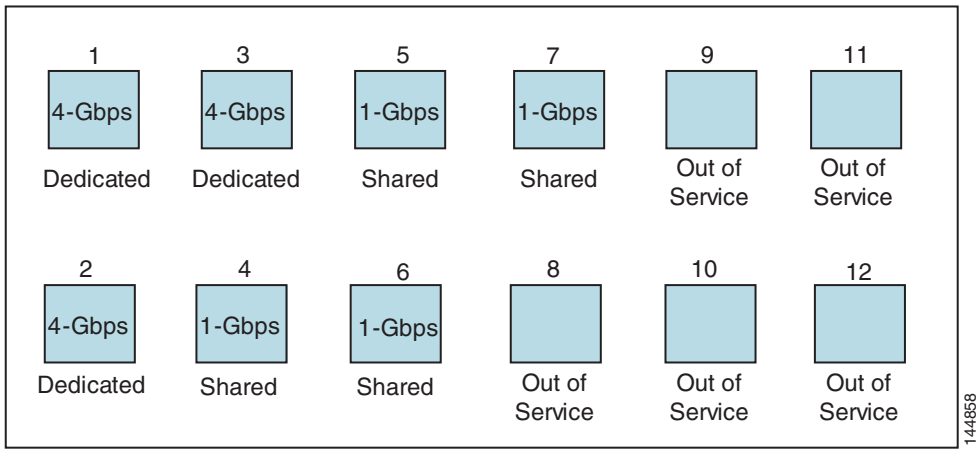

- Six ports with dedicated rate mode and 2-Gbps speed plus four ports with shared rate mode and 1-Gbps speed plus two ports put out-of-service (see Figure 19-3)

Figure 19-3 Example Speed and Rate Configuration on a 48-Port 4-Gbps Switching Module

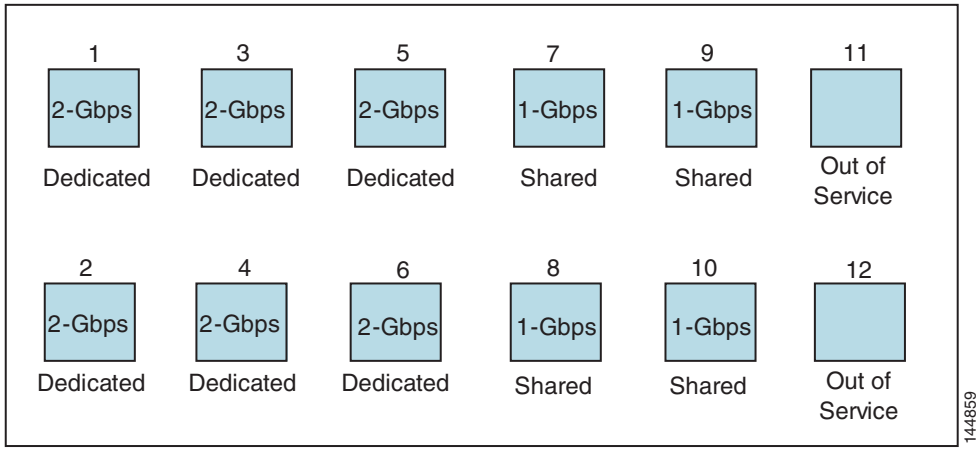

24-port 4-Gbps Switching Module BB_Credit Buffers

Table 19-5 lists the BB_credit buffer allocation for 24-port 4-Gbps switching modules.

Table 19-5 24-port 4-Gbps Switching Module BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Module	BB_Credit Buffers Per Port Defaults			
		Dedicated Rate Mode 4-Gbps Speed		Shared Rate Mode 4-Gbps Speed	
		ISL[1]	Fx Port	ISL[1]	Fx Port
User configurable BB_credit buffers	6000	250	16	16	16

1. ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 24-port 4-Gbps switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- Performance buffers are not supported on this module.

Each port group on the 24-port 4-Gbps switching module consists of six ports. The ports in shared rate mode have bandwidth oversubscription of 2:1 by default. However, some configurations of the shared ports in a port group can have maximum bandwidth oversubscription of 4:1 (considering that each port group has 12.8-Gbps bandwidth). The following example configurations are supported by the 24-port 4-Gbps switching modules:

- Six ports with shared rate mode and 4-Gbps speed (2:1 oversubscription) (default)
- Two ports with dedicated rate mode and 4-Gbps speed plus
  four ports with shared rate mode and 4-Gbps speed (with 4:1 oversubscription)
- One port with dedicated rate mode and 4-Gbps speed plus
  three ports with dedicated rate mode and 2-Gbps speed plus
  two ports with shared rate mode and 4-Gbps speed (4:1 oversubscription)
- Six ports with dedicated rate mode and 2-Gbps speed
- Three ports with dedicated rate mode and 4-Gbps speed plus
  three ports with shared rate mode and 1-Gbps speed (see Figure 19-4)

Figure 19-4 Example Speed and Rate Configuration on a 24-Port 4-Gbps Switching Module

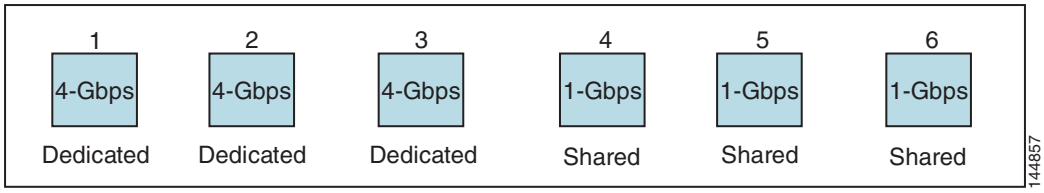

12-Port 4-Gbps Switching Module BB_Credit Buffers

Table 19-6 lists the BB_credit buffer allocation for 12-port 4-Gbps switching modules.

Table 19-6 12-Port 4-Gbps Switching Module BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Module	BB_Credit Buffers Per Port Defaults	
		Dedicated Rate Mode 4-Gbps Speed	
		ISL	Fx Port
User configurable BB_credit buffers	5488	250	16
Performance buffers	512 (shared)	145	12

The following considerations apply to BB_credit buffers on 12-port 4-Gbps switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers.
- BB_credit buffers for Fx port mode connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers.
- By default, 512 performance buffers are preallocated and are shared by all the ports. These buffers are configurable and the buffers are assigned to the port based on the availability of the buffers in the shared pool.
- There are 2488 extra buffers available as extended BB_credit buffers after allocating all the default BB_credit buffers for all the ports in ISL mode (5488 - (250 * 12)).

Note Extended BB_credits are allocated across all ports on the switch. That is, they are not allocated by port group.

Note By default, the ports in the 12-port 4-Gbps switching modules come up in 4-Gbps dedicated rate mode but can be configured as 1-Gbps and 2-Gbps dedicated rate mode. Shared mode is not supported.

4-Port 10-Gbps Switching Modules

Table 19-7 lists the BB_credit buffer allocation for 4-port 10-Gbps switching modules.

Table 19-7 4-port 10-Gbps Switching Module BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Module	BB_Credit Buffers Per Port Defaults	
		Dedicated Rate Mode 4-Gbps Speed	
		ISL[1]	F port[2]
User configurable BB_credit buffers	5488	250	16
Performance buffers	512 (shared)	145	12

1. ISL = E port or TE port.
2. Ports on the 4-port 10-Gbps cannot operate in FL port mode.

Note The ports in the 4-port 10-Gbps switching module only support 10-Gbps dedicated rate mode. FL port mode and shared rate mode are not supported.

The following considerations apply to BB_credit buffers on 4-port 10-Gbps switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers.
- BB_credit buffers for Fx port mode connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers.

- By default, 512 performance buffers are preallocated and are shared by all the ports. These buffers are configurable and the buffers are assigned to the port based on the availability of the buffers in the shared pool.
- There are 4488 extra buffers available as extended BB_credits after allocating all the default BB_credit buffers for all the ports in ISL mode (5488 - (250 * 4)).

Note Extended BB_credits are allocated across all ports on the switch. That is, they are not allocated by port group.

Extended BB_Credits

To facilitate BB_credits for long haul links, the extended BB_credits feature allows the user to configure the receive buffers above the maximum value on all Generation 2 switching modules (see the "Receive BB_Credit Buffers" section on page 19-5). When necessary, you can reduce the buffers on one port and assign them to another port, exceeding the default maximum. The minimum extended BB_credits per port is 256 and the maximum is 4095.In general, the user can configure any port in a port group to dedicated mode. To do this, you must first release the buffers from the other ports before configuring larger extended BB_credits for a port.

Note The ENTERPRISE_PKG license is required to use extended BB_credits on Generation 2 switching modules.

Note Extended BB_credits are not supported by ports in shared rate mode.

Note All ports on the Generation 2 switching modules support extended BB_credits. There are no limitations for how many extended BB_credits you can assign to a port (except for the maximum and minimum limits). If necessary, you can take interfaces out of service to make more extended BB_credits available to other ports.

Examples of Extended BB_Credit Configurations

Example 19-3 to Example 19-5 show extended BB_credit buffer configurations for each type of Generation 2 switching module.

Note The BB_credit values in the examples that are in bold font are extended buffer credit values.

Example 19-1 48-Port 4-Gbps Switching Module BB_Credit Configuration Example

	Port Numbers											
	1	2	3	4	5	6	7	8	9	10	11	12
BB_credits	**900**	16	16	16	16	16	16	16	16	16	16	16

	Port Numbers											
	13	14	15	16	17	18	19	20	21	22	23	24
BB_credits	**600**	4	**500**	4	10	10	10	10	16	16	16	16

	Port Numbers											
	25	26	27	28	29	30	31	32	33	34	35	36
BB_credits	**500**	**500**	16	16	16	10	10	10	10	14	14	14

	Port Numbers											
	37	38	39	40	41	42	43	44	45	46	47	48
BB_credits	**200**	**200**	**200**	**200**	125	125	125	125	125	125	125	125

Example 19-2 24-Port 4-Gbps Switching Module BB_Credit Configuration Example

	Port Numbers											
	1	2	3	4	5	6	7	8	9	10	11	12
BB_credits	**1000**	**800**	16	16	16	16	**500**	**200**	**200**	**200**	16	16

	Port Numbers											
	13	14	15	16	17	18	19	20	21	22	23	24
BB_credits	**1200**	**500**	**500**	**150**	**150**	**150**	16	16	16	16	16	16

Example 19-3 12-Port 4-Gbps Switching Module BB_Credit Configuration Example

	Port Numbers											
	1	2	3	4	5	6	7	8	9	10	11	12
BB_credits	**1500**	100	**1900**	100	**600**	**500**	50	250	**300**	100	**300**	200

Example 19-4 4-Port 10-Gbps Switching Module BB_Credit Configuration Example

	Port Numbers			
	1	2	3	4
BB_credits	1400	1000	**3000**	550

Example 19-5 4-Port 10-Gbps Switching Module BB_Credit Configuration Example

	Port Numbers			
	1	2	3	4
BB_credits	**4095**	250	200	32

About Combining Generation 1 Modules and Generation 2 Modules

All the existing Generation 1 and Generation 2 switching modules are supported by Cisco MDS SAN-OS Release 3.0(1) and later. However, there are limitations to consider when combining the various modules and supervisors in the Cisco MDS 9500 Series platform chassis.

This section includes the following topics:

Port Indexes

Cisco MDS 9000 switches allocate index identifiers for the ports on the modules. These port indexes cannot be configured. You can combine Generation 1 and Generation 2 switching modules, with either Supervisor-1 modules or Supervisor-2 modules. However, combining switching modules and supervisor modules has the following port index limitations:

- Supervisor-1 modules only support a maximum of 252 port indexes, regardless of the type of switching modules.
- Supervisor-2 modules support a maximum of 1020 port indexes when all switching modules in the chassis are Generation 2.
- Supervisor-2 modules only support a maximum of 252 port indexes when only Generation 1 switching modules, or a combination of Generation 1 and Generation 2 switching modules, are installed in the chassis.

Note On a switch with the maximum limit of 252 port index maximum limit, any new module that exceeds the limit when installed does not power up.

Generation 1 switching modules have specific numbering requirements. If these requirements are not met, the module does not power up. The port index numbering requirements include the following:

- If port indexes in the range of 256 to 1020 are assign to operational ports, Generation 1 switching modules do not power up.
- A block of contiguous port indexes is available. If such a block of port indexes is not available, Generation 1 modules do not power up. Table 19-8 shows the port index requirements for the Generation 1 modules.

Note If the switch has Supervisor-1 modules, the block of 32 contiguous port indexes must begin on the slot boundary. The slot boundary for slot 1 is 0, for slot 2 is 32, and so on. For Supervisor-2 modules, the contiguous block can start anywhere.

Table 19-8 Port Index Requirements for Generation 1 Modules

Generation 1 Module	Number of Port Indexes Required	
	Supervisor-1 Module	Supervisor-2 Module
16-port 2-Gbps Fibre Channel module	16	16
32-port 2-Gbps Fibre Channel module	32	32
8-port Gigabit Ethernet IP Storage Services module	32	32
4-port Gigabit Ethernet IP Storage Services module	32	16
32-port 2-Gbps Fibre Channel Storage Services Module (SSM).	32	32
14-port Fibre Channel/2-port Gigabit Ethernet Multiprotocol Services (MPS-14/2) module.	32	22

The allowed mix of Generation 1 and Generation 2 switching modules in a chassis is determined at run-time, either when booting up the switch or when installing the modules. In some cases, the sequence in which switching modules are inserted into the chassis determines if one or more modules is powered up. When a module does not power up because of a resource limitation, you can see the reason by viewing the module information in the Information pane.

The running configuration is updated when modules are installed. If you save the running configuration to the startup configuration during reboot, the switch powers up the same set of modules as before the reboot regardless of the sequence in which the modules initialize.

For information on recovering a module powered-down because port indexes are not available, refer to the *Cisco MDS 9000 Family Troubleshooting Guide, Release 3.x*.

PortChannels

PortChannels have the following restrictions:

- The maximum number of PortChannels allowed is 256 if all switching modules are Generation 1 or Generation 2.
- The maximum number of Port Channels allowed is 128 if the switching modules are both Generation 1 and Generation 2.

Note The number of PortChannels allowed does not depend on the type of supervisor module.

When configuring PortChannels on switches with both Generation 1 and Generation 2 modules, configure the PortChannel and Generation 2 switching modules interfaces to auto with a maximum of 2 Gbps or configure the Generation 1 switching modules followed by the Generation 2 switching modules.

Note Generation 1 switching module interfaces do not support auto speed with max 2Gbps. Also, Generation 2 switching module interfaces cannot be forcefully added to a PortChannel if sufficient resources are not available.

Note Before adding a Generation 2 interface to a PortChannel, check for resource availability.

Table 19-9 describes the results of adding a member to a PortChannel for various configurations.

Table 19-9 PortChannel Configuration and Addition Results

PortChannel Members	Configured Speed		New Member Type	Addition Type	Result
	PortChannel	Member			
No members	Any	Any	Generation 1 or Generation 2	Force	Pass
	Auto	Auto	Generation 1 or Generation 2	Normal or force	Pass
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 2	Normal	Fail
				Force	Fail
Generation 1 interfaces	Auto	Auto	Generation 2	Normal	Fail
				Force	Pass
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2	Normal	Fail
				Force	Pass
	Auto max 2000	Auto	Generation 2	Normal	Fail
				Force	Pass
Generation 2 interfaces	Auto	Auto	Generation 1	Normal or force	Fail
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 2	Normal	Fail
				Force	Pass

Configuring Generation 2 Module Interface Shared Resources

This section describes how to configure Generation 2 module interface shared resources and contains the following sections:

Configuration Guidelines for 24-Port and 48-Port 4-Gbps Switching Modules

The 24-port and 48-port 4-Gbps switching modules support the following features:

- 1-Gbps, 2-Gbps, and 4-Gbps speed traffic
- Shared and dedicated rate mode
- ISL (E or TE) and Fx (F or FL) port modes
- Extended BB_credits

Migrating from Shared Mode to Dedicated Mode

To configure 24-port and 48-port 4-Gbps switching modules when starting with the default configuration or when migrating from shared rate mode to dedicated rate mode, follow these guidelines:

1. Take unused interfaces out of service to release resources for other interfaces, if necessary.

See the "Taking Interfaces Out of Service" section on page 19-19.

2. Configure the traffic speed to use (1 Gbps, 2 Gbps, 4 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps).

See the "Configuring Port Speed" section on page 19-17.

3. Configure the rate mode (dedicated or shared) to use.

See the "Configuring Rate Mode" section on page 19-18.

4. Configure the port mode.

See the "About Interface Modes" section on page 18-3.

Note ISL ports cannot operate in shared rate mode.

5. Configure the BB_credits and extended BB_credits, as necessary.

See the "About Buffer-to-Buffer Credits" section on page 18-12 and the "About Extended BB_credits" section on page 18-13.

Migrating from Dedicated Mode to Shared Mode

To configure 24-port and 48-port 4-Gbps switching modules migrating from dedicated rate mode to shared rate mode, follow these guidelines:

1. Take unused interfaces out of service to release resources for other interfaces, if necessary.

See the "Taking Interfaces Out of Service" section on page 19-19.

2. Configure the BB_credits and extended BB_credits, as necessary.

See the "About Buffer-to-Buffer Credits" section on page 18-12 and the "About Buffer-to-Buffer Credits" section on page 18-12.

3. Configure the port mode.

See the "Configuring Interface Modes" section on page 18-6.

Note ISL ports cannot operate in shared rate mode.

4. Configure the rate mode (dedicated or shared) to use.

See the "Configuring Rate Mode" section on page 19-18.

5. Configure the traffic speed to use (1 Gbps, 2 Gbps, 4 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps).

See the "Configuring Port Speed" section on page 19-17.

Configuration Guidelines for 12-Port 4-Gbps Switching Module Interfaces

The 12-port 4-Gbps switching modules support the following features:

- 1-Gbps, 2-Gbps, and 4-Gbps speed traffic
- Only dedicated rate mode
- ISL (E or TE) and Fx (F or FL) port modes
- Extended BB_credits
- Performance buffers

To configure 4-port 10-Gbps switching modules when starting with the default configuration, follow these guidelines:

1. Configure the traffic speed to use (1 Gbps, 2 Gbps, 4 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps).

See the "Configuring Port Speed" section on page 19-17.

2. Configure the port mode.

See the "Configuring Interface Modes" section on page 18-6.

3. Configure the BB_credits, performance buffers, and extended BB_credits, as necessary.

See the "About Buffer-to-Buffer Credits" section on page 18-12, the "About Performance Buffers" section on page 18-13, and the "About Buffer-to-Buffer Credits" section on page 18-12.

Note If you change the port bandwidth reservation parameters on a 24- or 48-port module, the change affects only the changed port. No other ports in the port group are affected.

Configuration Guidelines for 4-Port 10-Gbps Switching Module Interfaces

The 4-port 10-Gbps switching modules support the following features:

- Only 10-Gbps speed traffic
- Only dedicated rate mode
- ISL (E or TE) and F port modes
- Extended BB_credits
- Performance buffers

Use the following guidelines to configure 4-port 10-Gbps switching modules when starting with the default configuration:

1. Configure the port mode.

   See the "About Interface Modes" section on page 18-3.

2. Configure the BB_credits, performance buffers, and extended BB_credits, as necessary.

See the "About Buffer-to-Buffer Credits" section on page 18-12, the "About Performance Buffers" section on page 18-13, and the "About Buffer-to-Buffer Credits" section on page 18-12.

Configuring Port Speed

The port speed on an interface, combined with the rate mode, determines the amount of shared resources available to the ports in the port group on a 24-port or 48-port 4-Gbps switching module. Especially in the case of dedicated rate mode, the port group resources are reserved even though the bandwidth is not used. For example, if an interface is configured for autosensing (**auto**), then 4 Gbps of bandwidth is reserved even though the maximum operating speed is 2 Gbps. For the same interface, if autosensing with a maximum speed of 2 Gbps (**auto max 2000**) is configured, then only 2 Gbps of bandwidth is reserved and the unused 2 Gbps is shared with the other interface in the port group.

Note The 4-port 10-Gbps switching module only supports 10-Gbps traffic.

To configure dedicated bandwidth on an interface using Fabric Manager, follow these steps:

Step 1 Select a switch from the Fabric pane, or select a group of switches (SAN, fabric, VSAN) from the Logical Domains pane.

Step 2 Expand **Switches,** expand **Interfaces** and select **FC Physical** from the Physical Attributes pane.

You see the **FC Physical > General** tab in the Interfaces pane.

Step 3 Scroll until you see the row containing the switch and port you want to configure.

Step 4 Select **auto**, **1Gb**, **4Gb**, or **autoMax2G** from the Speed Admin column (see Figure 19-5).

Figure 19-5 Speed Admin column in Port Configuration

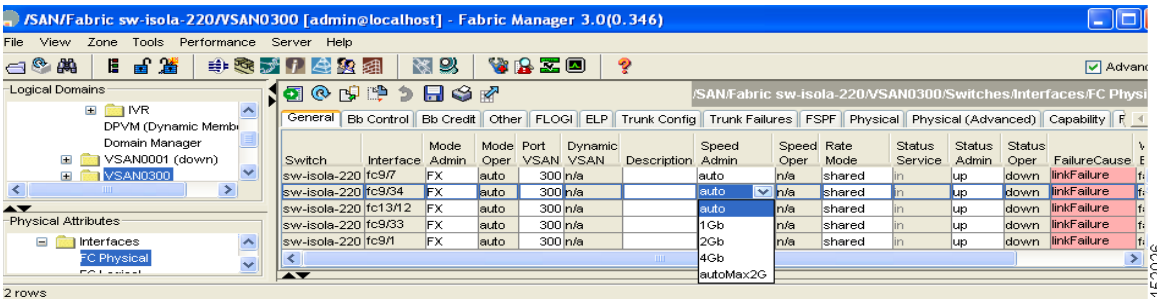

The **auto** parameter enables autosensing on the interface. The **autoMax2G** parameter enables autosensing on the interface with a maximum speed of 2 Gbps.

Note If you change the port bandwidth reservation parameters on a 24- or 48-port module, the change affects only the changed port. No other ports in the port group are affected.

Step 5 Click the **Apply Changes** icon.

Configuring Rate Mode

To configure the rate mode (dedicated or shared) on an interface on a 24-port 4-Gbps switching module or 48-port 4-Gbps switching module, follow these steps:

Step 1 Select a switch from the Fabric pane, or select a group of switches (SAN, fabric, VSAN) from the Logical Domains pane.

Step 2 Expand **Switches > Interfaces** and then select **FC Physical** from the Physical Attributes pane.

You see the **FC Physical > General** tab in the Interfaces pane.

Step 3 Scroll until you see the row containing the switch and port you want to configure.

Step 4 Select **dedicated** or **shared** from the Rate Mode column (see Figure 19-6).

Figure 19-6 Rate Mode Port Configuration

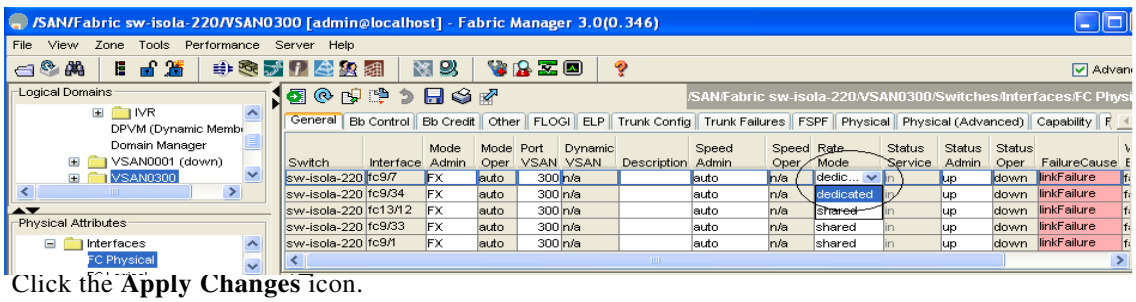

Step 5 Click the **Apply Changes** icon.

Taking Interfaces Out of Service

You can take interfaces out of service on Generation 2 switching modules. When an interface is out of service, all the shared resources for the interface are released as well as the configuration associated with those resources.

Note The interface must be disabled before it can be taken out of service.

Caution Taking interfaces out of service releases all the shared resources to ensure that they are available to other interfaces. This causes the configuration in the shared resources to revert to default when the interface is brought back into service. Also, an interface cannot come back into service unless the default shared resources for the port are available. The operation to free up shared resources from another port is disruptive.

Note The interface cannot be a member of a PortChannel.

When an interface is out of service, all the shared resources for the interface are released as well as the configuration associated with those resources.

Note The interface must be disabled before it can be taken out of service.

To take an interface out of service using Fabric Manager, follow these steps:

Step 1 Select a switch from the Fabric pane, or select a group of switches (SAN, fabric, VSAN) from the Logical Domains pane.

Step 2 Expand **Switches**, expand **Interfaces** and select **FC Physical** in the Physical Attributes pane.

You see the **FC Physical > General** tab in the Information pane.

Step 3 Scroll down until you see the row containing the switch and port you want to configure.

Step 4 Scroll right (if necessary) until you see the Status Service column.

Step 5 Select **in** or **out** from the Status Service column.

Step 6 Click the **Apply Changes** icon.

Releasing Shared Resources in a Port Group

When you want to reconfigure the interfaces in a port group on a Generation 2 module, you can return the port group to the default configuration to avoid problems with allocating shared resources.

Note The interface cannot be a member of a PortChannel.

Caution Releasing shared resources is a disruptive operation.

To release the shared resources for a port group using Fabric Manager, follow these steps:

Step 1 Select a switch from the Fabric pane, or select a group of switches (SAN, fabric, VSAN) from the Logical Domains pane.

Step 2 Expand **Switches > Interfaces** and then select **FC Physical** from the Physical Attributes pane.

You see the **FC Physical > General** tab in the Information pane.

Step 3 Scroll down until you see the row containing the switch and port you want to configure.

Step 4 Scroll right (if necessary) until you see the Status Service column (see Figure 19-7).

Figure 19-7 Status Service Column for FC Physical

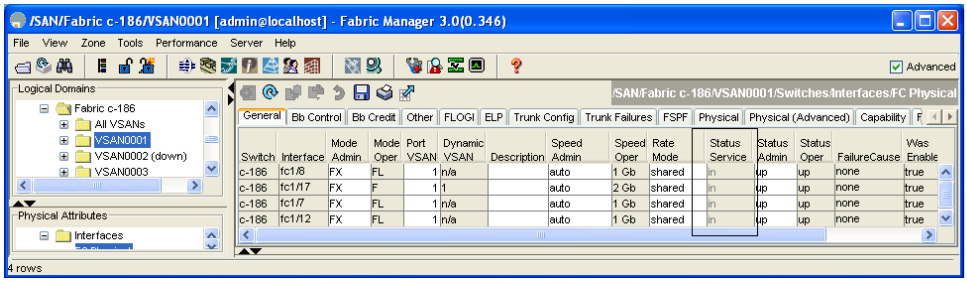

Step 5 Select the **out** status from the Status Service column.

Step 6 Click the **Apply Changes** icon.

Step 7 Select the **in** status from the Status Service column.

Step 8 Click the **Apply Changes** icon.

Configuring the Buffer-to-Buffer State Change Number

The BB_SC_N field (word 1, bits 15-12) specifies the Buffer-to-Buffer state change number. It indicates that the sender of the port login (PLOGI) or fabric login (FLOGI) frame is requesting 2BB_SC_N number of frames to be sent between two consecutive BB_SCs primitives, and 2BB_SC_N number of R_RDY primitives to be sent between two consecutive BB_SCr primitives.

To use the BB_SC_N field during PLOGI or FLOGI, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1` `switch(config-if)#`	Selects the interface and enters interface configuration submode.

	Command	Purpose
Step 3	`switch(config-if)# switchport fcbbscn`	Enables the use of buffer-to-buffer state change number for PLOGIs and FLOGIs on the interface.
	`switch(config-if)# no switchport fcbbscn`	Disables (default) the use of buffer-to-buffer state change number for PLOGIs and FLOGIs on the interface.

Default Settings

Table 19-10 lists the default settings for Generation 2 interface parameters.

Table 19-10 Default Generation 2 Interface Parameters

Parameter	Default			
	48-Port 4-Gbps Switching Module	24-Port 4-Gbps Switching Module	12-Port 4-Gbps Switching Module	4-Port 10-Gbps Switching Module
Speed mode	**auto**[1]	**auto**[1]	**auto**[1]	**auto**[2]
Rate mode	**shared**	**shared**	**dedicated**	**dedicated**
Port mode	**Fx**	**Fx**	**auto**[3]	**auto**[4]
BB_credit buffers	16	16	250	250
Performance buffers	–	–	145[5]	145[5]

1. Auto speed mode on the 4-Gbps switching modules negotiates 1, 2, and 4 Gbps.
2. The 4-port 10-Gbps switching module only supports 10-Gbps traffic.
3. Auto port mode on the 12-port 4-Gbps switching module interfaces can operate in E port mode, TE port mode, and Fx port mode.
4. Auto port mode on the 4-port 10-Gbps switching module interfaces can operate in E port mode, TE port mode, and F port mode.
5. Performance buffers are shared among all ports on the module.

CHAPTER 20

Configuring Trunking

This chapter describes the trunking feature provided in Cisco MDS 9000 switches. It includes the following sections:

About Trunking

Trunking, also known as VSAN trunking, is a feature specific to switches in the Cisco MDS 9000 Family. Trunking enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link, using Enhanced ISL (EISL) frame format (see Figure 20-1).

Figure 20-1 Trunking

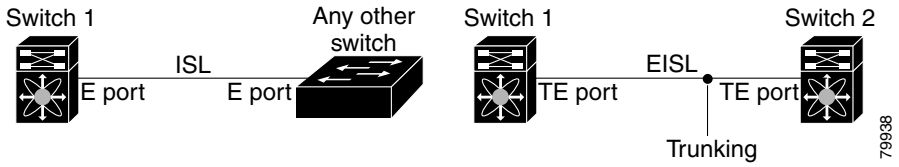

The trunking feature includes the following restrictions:

- Trunking configurations are only applicable to E ports. If trunk mode is enabled in an E port and that port becomes operational as a trunking E port, it is referred to as a TE port.
- The trunk-allowed VSANs configured for TE ports are used by the trunking protocol to determine the allowed-active VSANs in which frames can be received or transmitted.
- If a trunking enabled E port is connected to a third-party switch, the trunking protocol ensures seamless operation as an E port.

If you misconfigure VSAN configurations across E ports, you could face consequences such as merging the traffic in two VSANs (thus causing both VSANs to mismatch). The trunking protocol validates the VSAN interfaces at both ends of an ISL to avoid merging VSANs (see Figure 20-2).

Figure 20-2 VSAN Mismatch

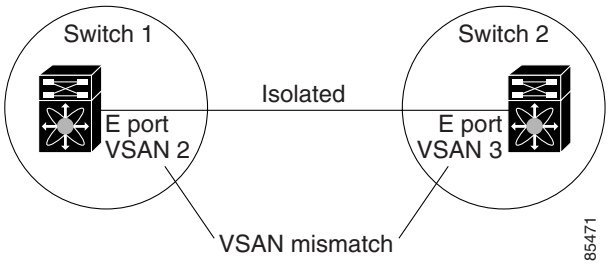

In this example, the trunking protocol detects potential VSAN merging and isolates the ports involved.

The trunking protocol cannot detect merging of VSANs when a third-party switch is placed in between two Cisco MDS 9000 Family switches (see Figure 20-3).

Figure 20-3 Third-Party Switch VSAN Mismatch

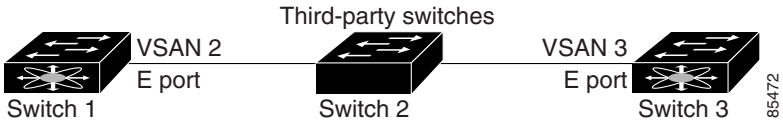

VSANs 2 and 3 get effectively merged with overlapping entries in the name server and the zone applications.

Trunking Protocol

The trunking protocol is important for E-port and TE-port operations. It supports the following:

- Dynamic negotiation of operational trunk mode.
- Selection of a common set of trunk-allowed VSANs.
- Detection of a VSAN mismatch across an ISL.

By default, the trunking protocol is enabled. If the trunking protocol is disabled on a switch, no port on that switch can apply new trunk configurations. Existing trunk configurations are not affected—the TE port continues to function in trunk mode, but only supports traffic in VSANs that it negotiated with previously (when the trunking protocol was enabled). Also, other switches that are directly connected to this switch are similarly affected on the connected interfaces. In some cases, you may need to merge traffic from different port VSANs across a non-trunking ISL. If so, disable the trunking protocol.

Tip To avoid inconsistent configurations, shut all E ports before enabling or disabling the trunking protocol.

This section explains how to configure trunking and contains the following topics:

About Trunk Mode

By default, trunk mode is enabled in all Fibre Channel interfaces. However, trunk mode configuration takes effect only in E-port mode. You can configure trunk mode as on (enabled), off (disabled), or auto (automatic). The default trunk mode is on. The trunk mode configuration at the two ends of an ISL, between two switches, determine the trunking state of the link and the port modes at both ends (see Table 20-1).

Table 20-1 Trunk Mode Status Between Switches

Your Trunk Mode Configuration		Resulting State and Port Mode	
Switch 1	**Switch 2**	**Trunking State**	**Port Mode**
On	Auto or on	Trunking (EISL)	TE port
Off	Auto, on, or off	No trunking (ISL)	E port
Auto	Auto	No trunking (ISL)	E port

Tip The preferred configuration on the Cisco MDS 9000 Family switches is one side of the trunk set to auto and the other set to on.

Note When connected to a third-party switch, the trunk mode configuration has no effect—the ISL is always in a trunking disabled state.

Figure 20-4 shows an example of configuring trunk mode using Fabric Manager.

Figure 20-4 Configuring Trunk Mode in Fabric Manager.

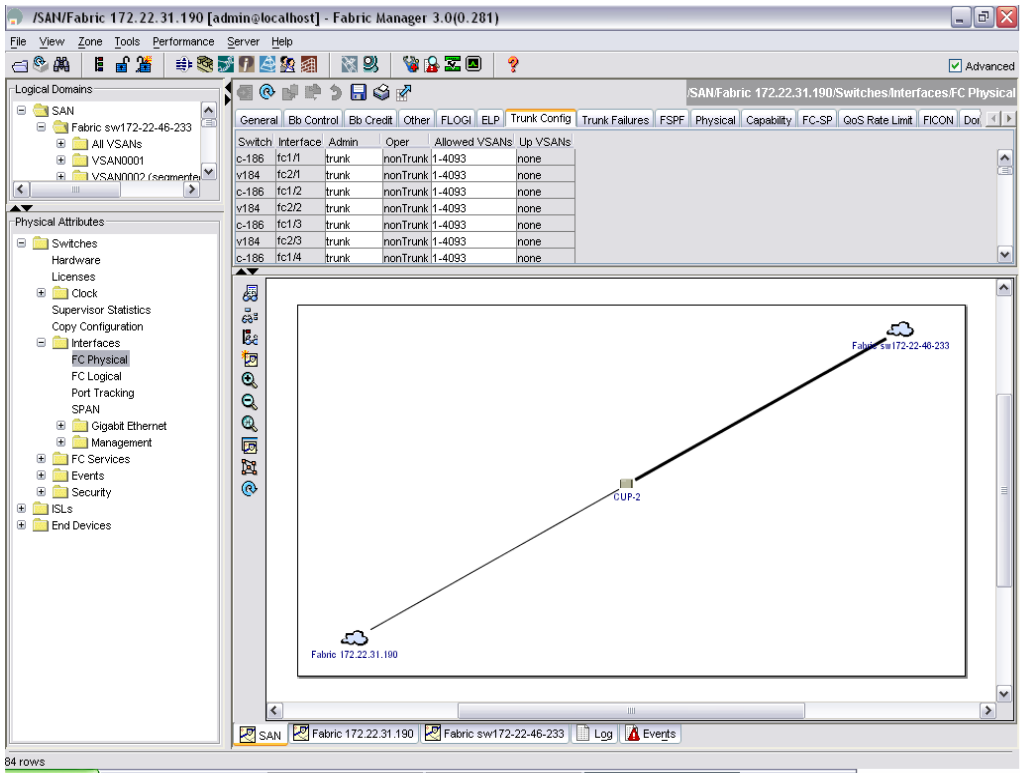

Configuring Trunk Mode

To configure trunk mode using Fabric Manager, follow these steps:

Step 1 Expand **Interfaces** and then select **FC Physical**. You see the interface configuration in the Information pane.

Step 2 Click the **Trunk Config** tab to modify the trunking mode for the selected interface.

You see the information shown in Figure 20-5.

Figure 20-5 Trunking Configuration

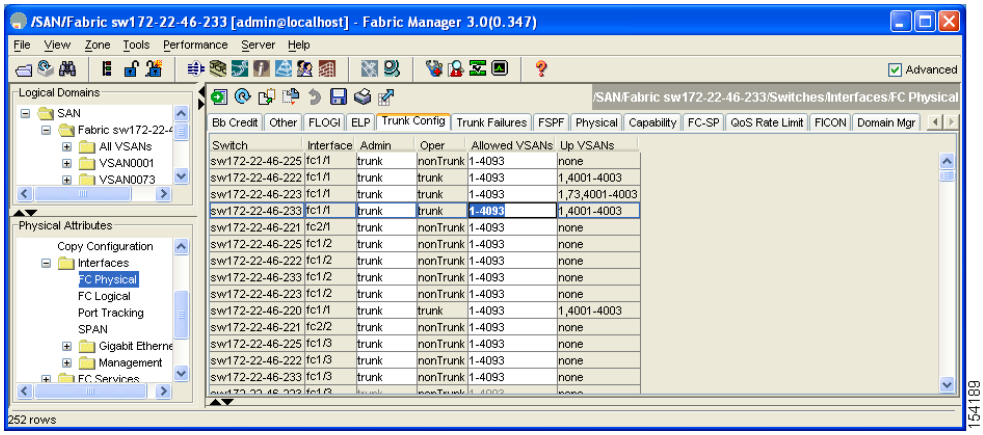

Step 3 Make changes to the Admin and Allowed VSANs values.

Step 4 Click **Apply Changes** or click **Undo Changes**.

About Allowed-Active VSAN Lists

Each Fibre Channel interface has an associated trunk-allowed VSAN list. In TE-port mode, frames are transmitted and received in one or more VSANs specified in this list. By default, the VSAN range (1 through 4093) is included in the trunk-allowed list.

The common set of VSANs that are configured and active in the switch are included in the trunk-allowed VSAN list for an interface, and they are called *allowed-active* VSANs. The trunking protocol uses the list of allowed-active VSANs at the two ends of an ISL to determine the list of operational VSANs in which traffic is allowed.

In Figure 20-6, switch 1 has VSANs 1 through 5, switch 2 has VSANs 1 through 3, and switch 3 has VSANs 1, 2, 4, and 5 with a default configuration of trunk-allowed VSANs. All VSANs configured in all three switches are allowed-active. However, only the common set of allowed-active VSANs at the ends of the ISL become operational as shown in Figure 20-6.

Figure 20-6 Default Allowed-Active VSAN Configuration

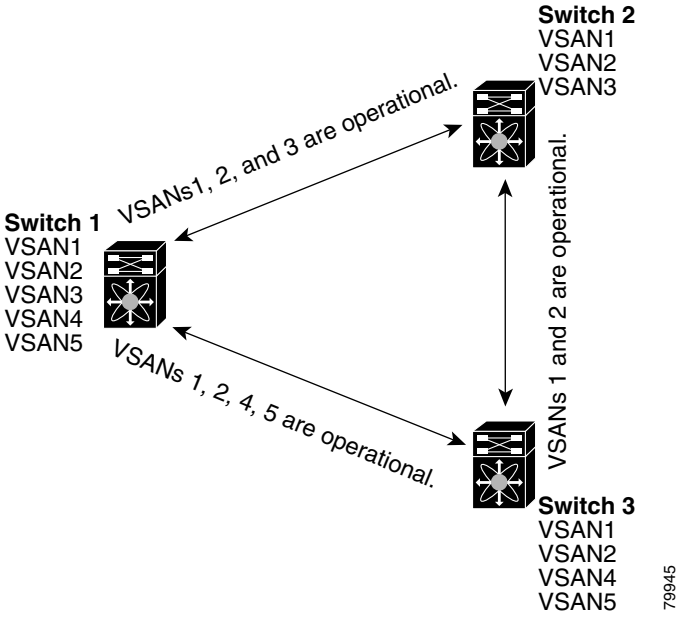

You can configure a select set of VSANs (from the allowed-active list) to control access to the VSANs specified in a trunking ISL.

Using Figure 20-6 as an example, you can configure the list of allowed VSANs on a per-interface basis (see Figure 20-7). For example, if VSANs 2 and 4 are removed from the allowed VSAN list of ISLs connecting to switch 1, the operational allowed list of VSANs for each ISL would be as follows:

- The ISL between switch 1 and switch 2 shall include VSAN 1 and VSAN 3.
- The ISL between switch 2 and switch 3 shall include VSAN 1 and VSAN 2.
- The ISL between switch 3 and switch 1 shall include VSAN 1, 2, and 5.

Consequently, VSAN 2 can only be routed from switch 1 through switch 3 to switch 2.

Figure 20-7 Operational and Allowed VSAN Configuration

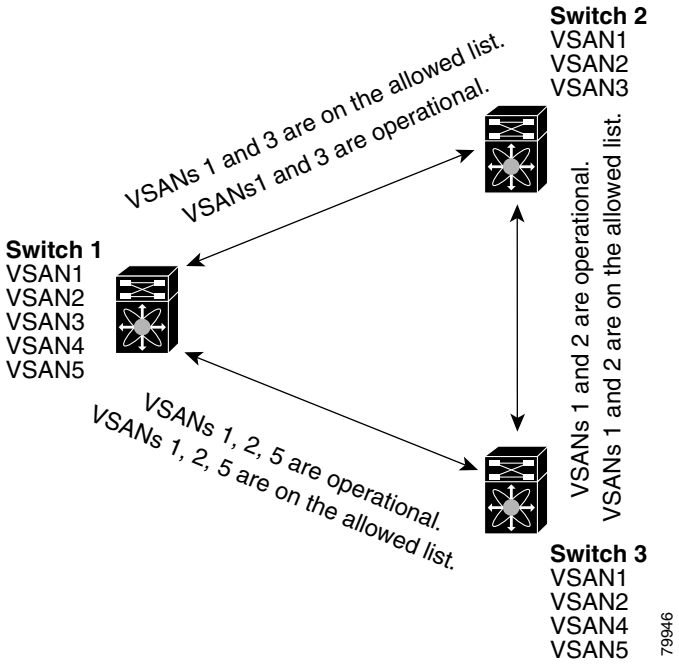

Configuring an Allowed-Active List of VSANs

To configure an allowed-active list of VSANs for an interface using Fabric Manager, follow these steps:

Step 1 Expand **Interfaces** and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **Trunk Config** tab.

You see the current trunk configuration.

Step 3 Set Allowed VSANs to the list of allowed VSANs for each interface that you want to configure.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

Trunking Configuration Guidelines

If you misconfigure VSAN configurations across E ports, you could face consequences such as merging the traffic in two VSANs (thus causing both VSANs to mismatch). The trunking protocol validates the VSAN interfaces at both ends of an ISL to avoid VSANs merging (see Figure 20-2).

Figure 20-8 VSAN Mismatch

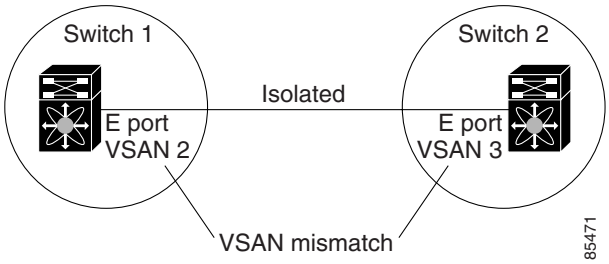

In this example, the trunking protocol detects potential VSAN merging and isolates the ports involved.

The trunking protocol cannot detect merging of VSANs when a third-party switch is placed in between two Cisco MDS 9000 Family switches (see Figure 20-3).

Figure 20-9 Third-Party Switch VSAN Mismatch

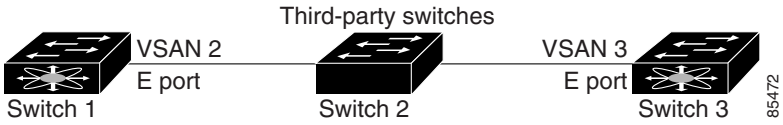

VSAN 2 and VAN 3 are effectively merged with overlapping entries in the name server and the zone applications. The Cisco MDS 9000 Fabric Manager helps detect such topologies.

Default Settings

Table 20-2 lists the default settings for trunking parameters.

Table 20-2 Default Trunk Configuration Parameters

Parameters	Default
Switch port trunk mode	On.
Allowed VSAN list	1 to 4093 user-defined VSAN IDs.
Trunking protocol	Enabled.

CHAPTER 21

Configuring PortChannels

PortChannels refer to the aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy. PortChannels can connect to interfaces across switching modules, so a failure of a switching module cannot bring down the PortChannel link.

This chapter discusses the PortChannel feature provided in the switch and includes the following sections:

About PortChannels

A PortChannel has the following functionality:

- Provides a point-to-point connection over ISL (E ports) or EISL (TE ports). Multiple links can be combined into a PortChannel.
- Increases the aggregate bandwidth on an ISL by distributing traffic among all functional links in the channel.
- Load balances across multiple links and maintains optimum bandwidth utilization. Load balancing is based on the source ID, destination ID, and exchange ID (OX ID).
- Provides high availability on an ISL. If one link fails, traffic previously carried on this link is switched to the remaining links. If a link goes down in a PortChannel, the upper protocol is not aware of it. To the upper protocol, the link is still there, although the bandwidth is diminished. The routing tables are not affected by link failure. PortChannels may contain up to 16 physical links and may span multiple modules for added high availability.

Note See the "Fail-Over Scenarios for PortChannels and FSPF Links" section on page 28-3 for fail-over scenarios.

Cisco MDS 9000 Family of switches support 128 PortChannels with 16 interfaces per PortChannel. A PortChannel number refers to the unique (to each switch) identifier associated with each channel group. This number ranges from 1 to 128.

This section describes PortChannels and contains the following topics:

PortChannel Examples

PortChannels on Cisco MDS 9000 Family switches allow flexibility in configuration. Figure 21-1 illustrates three possible PortChannel configurations:

- PortChannel A aggregates two links on two interfaces on the same switching module at each end of a connection.
- PortChannel B also aggregates two links, but each link is connected to a different switching module. If the switching module goes down, traffic is not affected.
- PortChannel C aggregates three links. Two links are on the same switching module at each end, while one is connected to a different switching module on switch 2.

Figure 21-1 PortChannel Flexibility

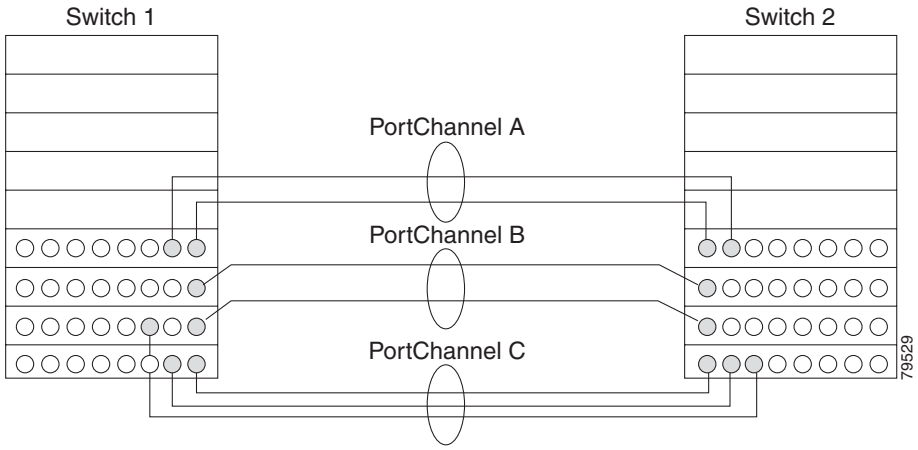

32-Port Switching Module Configuration Guidelines

The 32-port switching module guidelines applies to the following hardware:

- The 32-port 2-Gbps or 1-Gbps switching module
- The Cisco MDS 9140 Switch

When configuring these host-optimized ports, the following PortChannel guidelines apply:

- Any (or all) full line rate port(s) in the Cisco MDS 9100 Series can be included in a PortChannel.
- The host-optimized ports in the Cisco MDS 9100 Series are subject to the same PortChannel rules as 32-port switching modules—only the first port of each group of 4 ports is included in a PortChannel.
  - You can configure only the first port in each 4-port group (for example, the first port in ports 1-4, the fifth port in ports 5-8 and so on) as an E port. If the first port in the group is configured as a PortChannel, the other three ports in each group (ports 2-4, 6-8 and so on) are not usable and remain in the shutdown state.
  - If any of the other three ports are configured in a no shutdown state, you cannot configure the first port to be a PortChannel. The other three ports continue to remain in a no shutdown state.

Note In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the Cisco MDS 9120 Switch and 8 ports in the Cisco MDS 9140 Switch) are full line rate like the 16-port switching module. The other ports (16 ports in the Cisco MDS 9120 Switch and 32 ports in the Cisco MDS 9140 Switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.

About PortChanneling and Trunking

Trunking is a commonly used storage industry term. However, the Cisco SAN-OS software and switches in the Cisco MDS 9000 Family implement trunking and PortChanneling follows:

- PortChanneling enables several physical links to be combined into one aggregated logical link.
- Trunking enables a link transmitting frames in the EISL format to carry (trunk) multiple VSAN traffic. When trunking is operational on an E port, that E port becomes a TE port. A TE port is specific to switches in the Cisco MDS 9000 Family. An industry standard E port can link to other vendor switches and is referred to as a nontrunking interface (see Figure 21-2 and Figure 21-3).

See Chapter 20, "Configuring Trunking" for information on trunked interfaces.

Figure 21-2 Trunking Only

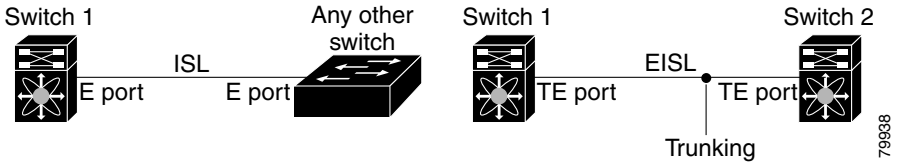

Figure 21-3 PortChanneling and Trunking

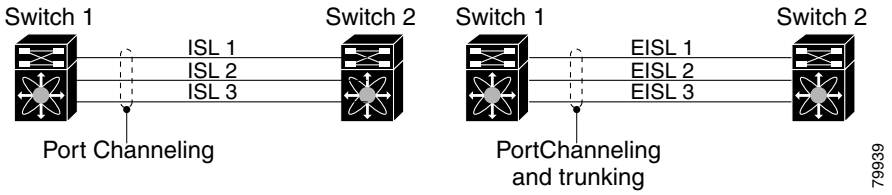

PortChanneling and trunking are used separately across an ISL:

- PortChanneling—Interfaces can be channeled between E ports and TE ports.
- Trunking—Trunking permits carrying traffic on multiple VSANs between switches. Interfaces can be trunked only between TE ports.

  See Chapter 23, "Configuring and Managing VSANs."

Both PortChanneling and trunking can be used between TE ports over EISLs.

About Load Balancing

Two mechanisms support the load balancing functionality:

- Flow based—All frames between source and destination follow the same links for a given flow. That is, whichever link is selected for the first exchange of the flow is used for all subsequent exchanges.
- Exchange based—The first frame in an exchange picks a link and subsequent frames in the exchange follow the same link. However, subsequent exchanges can use a different link. This provides more granular load balancing while preserving the order of frames for each exchange.

Figure 21-4 illustrates how source ID 1 (SID1) and destination ID1 based(DID1) load balancing works. When the first frame in a flow is received on an interface for forwarding, link 1 is selected. Each subsequent frame in that flow is sent over the same link. No frame in SID1 and DID1 utilizes link 2.

Figure 21-4 SID1 and DID1 Based Load Balancing

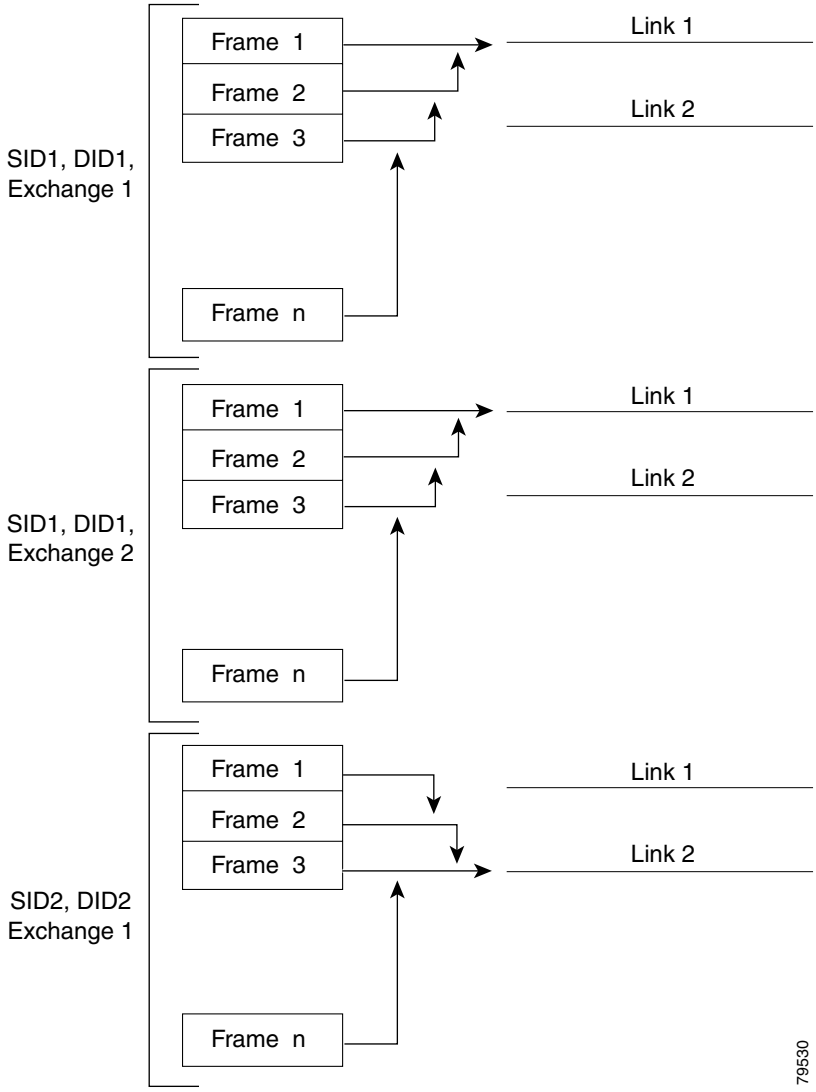

Figure 21-5 illustrates how exchange based load balancing works. When the first frame in an exchange is received for forwarding on an interface, link 1 is chosen by a hash algorithm. All remaining frames in that particular exchange are sent on the same link. For exchange 1, no frame uses link 2. For the next exchange, link 2 is chosen by the hash algorithm. Now all frames in exchange 2 use link 2.

Figure 21-5 SID1, DID1, and Exchange Based Load Balancing

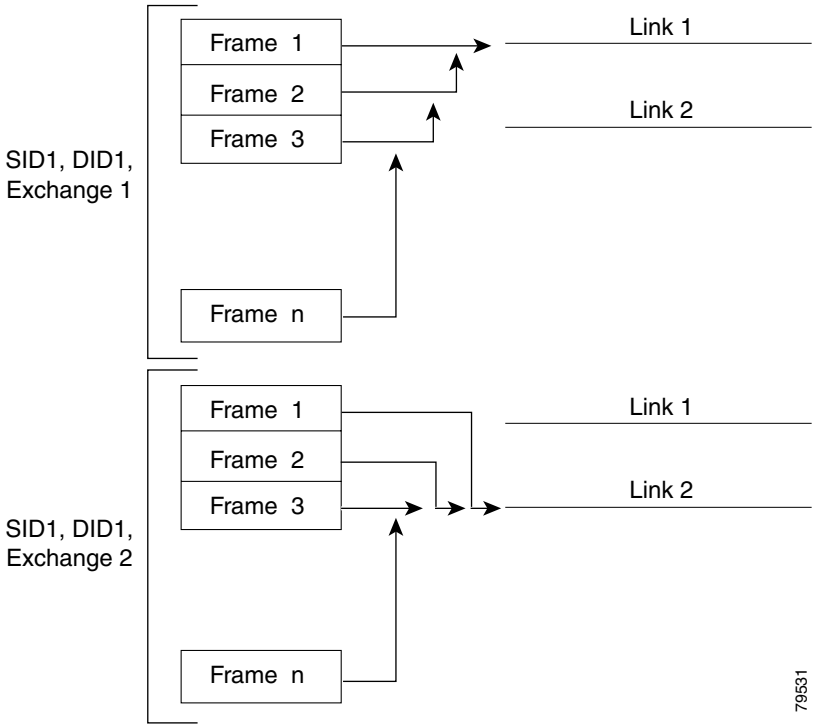

For more information on configuring load balancing and in-order delivery features, see the "About VSANs" section on page 23-1.

PortChannel Configuration

PortChannels are created with default values. You can change the default configuration just like any other physical interface.

Figure 21-6 provides examples of valid PortChannel configurations.

Figure 21-6 Valid Configurations

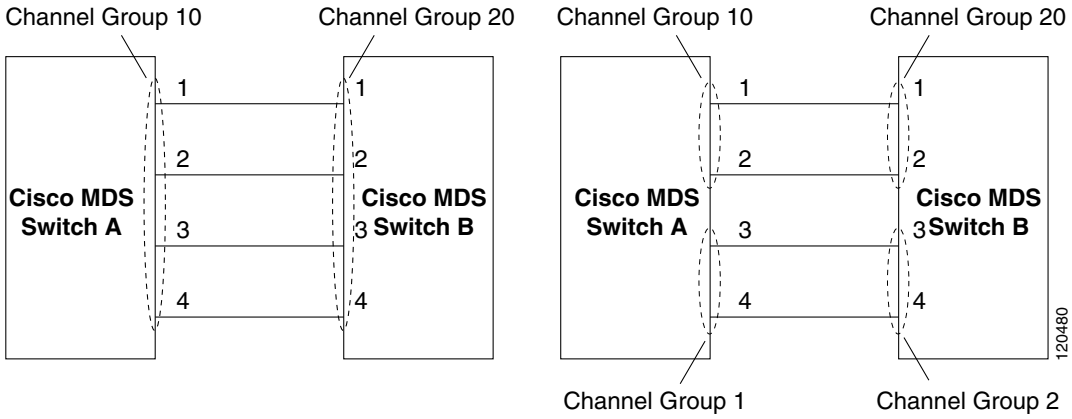

Figure 21-7 provides examples of invalid configurations. Assuming that the links are brought up in the 1,2,3, 4 sequence, links 3 and 4 will be operationally down as the fabric is misconfigured.

Figure 21-7 Misconfigured Configurations

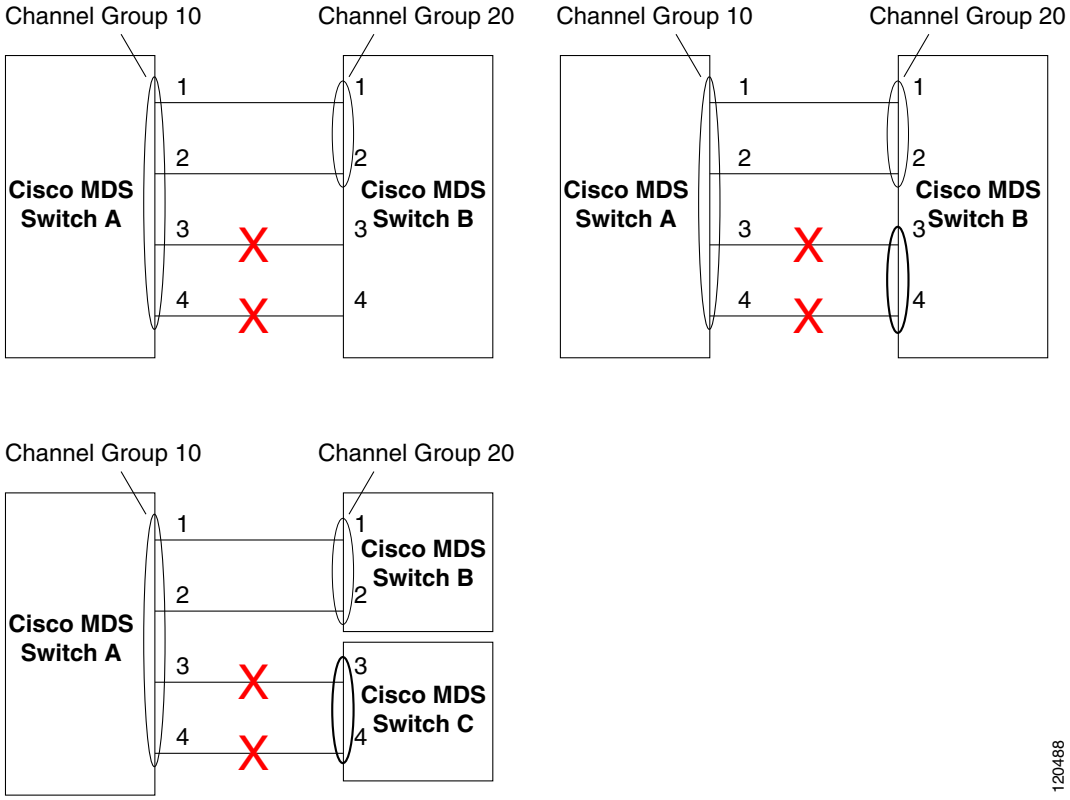

This section shows how to configure and modify PortChannels and contains the following topics:

About PortChannel Configuration

Before configuring a PortChannel, consider the following guidelines:

- Configure the PortChannel across switching modules to prevent redundancy on switching module reboots or upgrades.
- Ensure that one PortChannel is not connected to different sets of switches. PortChannels require point-to-point connections between the same set of switches.

If you misconfigure PortChannels, you may receive a misconfiguration message. If you receive this message, the PortChannel's physical links are disabled because an error has been detected.

A PortChannel error is detected if the following requirements are not met:

- Each switch on either side of a PortChannel must be connected to the same number of interfaces.
- Each interface must be connected to a corresponding interface on the other side (see Figure 21-7 for an example of an invalid configuration).
- Links in a PortChannel cannot be changed after the PortChannel is configured. If you change the links after the PortChannel is configured, be sure to reconnect the links to interfaces within the PortChannel and reenable the links.

If all three conditions are not met, the faulty link is disabled.

Configuring PortChannels

To create a PortChannel using the PortChannel Wizard in Fabric Manager, follow these steps:

Step 1 Click the **PortChannel Wizard** icon in the toolbar (see Figure 21-8).

Figure 21-8 PortChannel Wizard Icon

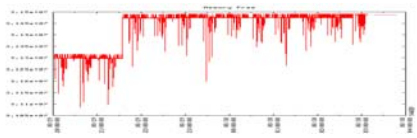

You see the first PortChannel Wizard screen (see Figure 21-9).

Step 2 Select a switch pair. Figure 21-9 shows a list of the switch pairs.

Figure 21-9 Select Switch Pairs

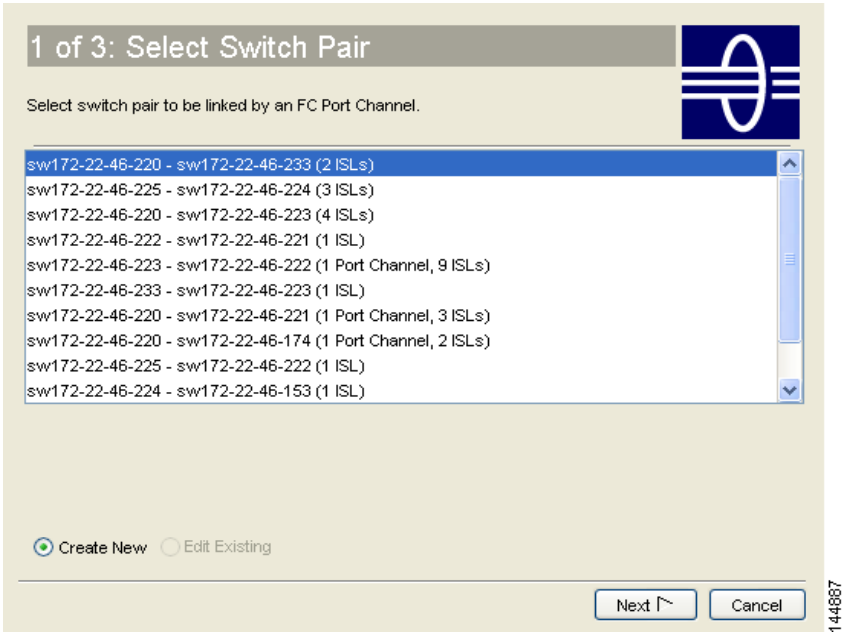

Step 3 Click **Next**.

Step 4 Select the ISLs. Figure 21-10 shows a list of the ISLs.

Figure 21-10 Select ISLs

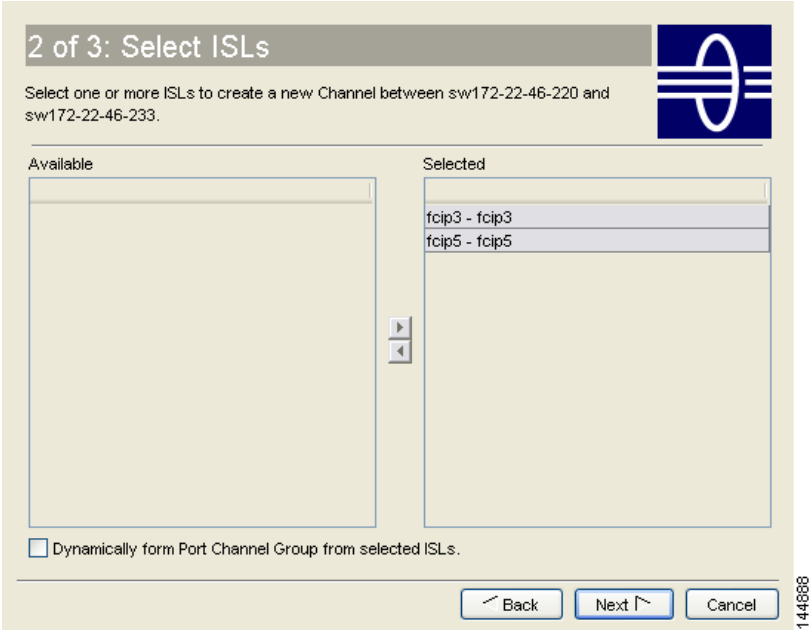

Step 5 Optionally, check the **Dynamically form Port Channel Group from selected ISLs** check box if you want to dynamically create the PortChannel and make the ISL properties identical for the Admin, Trunk, Speed, and VSAN attributes.

Step 6 Click **Next**.

Step 7 If you chose to dynamically form a PortChannel from selected ISLs, you see the final PortChannel Wizard screen as shown in Figure 21-11. Set the VSAN List, Trunk Mode, and Speed and proceed to Step 11.

Figure 21-11 Dynamically Form a PortChannel

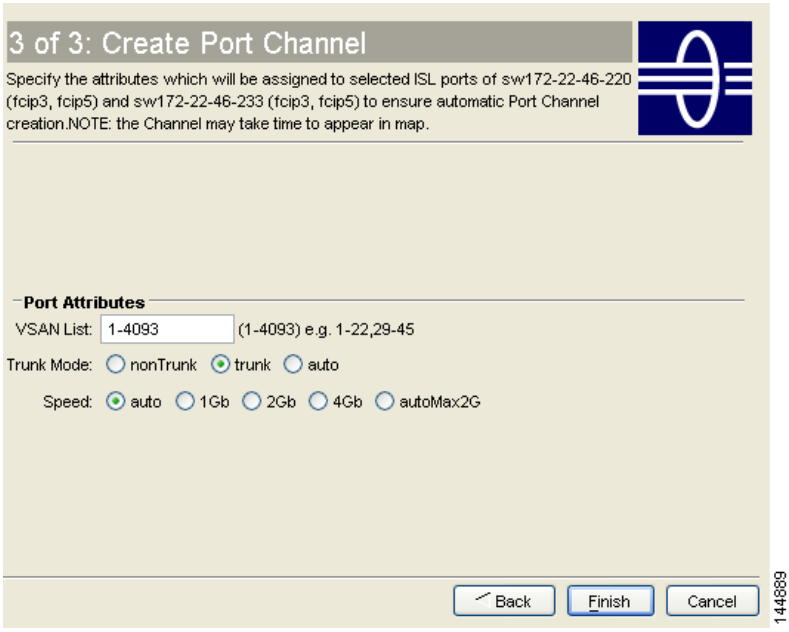

Step 8 If you did not choose to dynamically form a PortChannel, you see the third PortChannel Wizard screen as shown in Figure 21-12.

Figure 21-12 Create a PortChannel

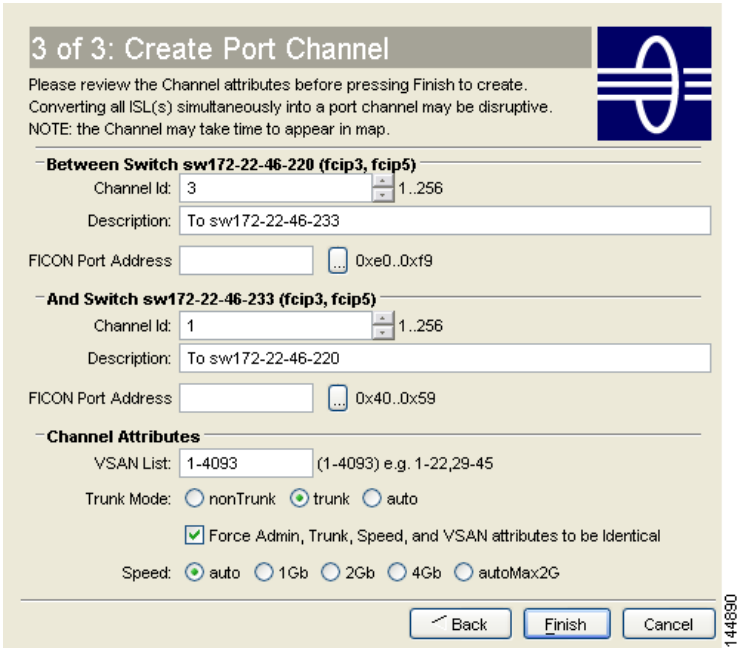

Step 9 Change the channel ID or description for each switch, if necessary.

Step 10 Review the attributes at the bottom of the screen, and set them if applicable.

The following attributes are shown in Figure 21-12:

- VSAN List—This gives a list of VSANs to which the ISLs belong.
- Trunk Mode—You can enable trunking on the links in the PortChannel. Select **trunking** if your link is between TE ports. Select **nontrunking** if your link is between E ports (for example, if your link is between an MDS switch and another vendor's switch). Select **auto** if you are not sure.
- Force Admin, Trunk, Speed, and VSAN attributes to be identical—This check box ensures that the same parameter settings are used in all physical ports in the channel. If these settings are not identical, the ports cannot become part of the PortChannel.
- Speed—The port speed values are **auto**, **1Gb**, **2Gb**, **4Gb**, and **autoMax2G**.

Step 11 Click **OK**.

The PortChannel is created. Note that it may take a few minutes before the new PortChannel is visible in the Fabric pane.

About PortChannel Modes

You can configure each PortChannel with a channel group mode parameter to determine the PortChannel protocol behavior for all member ports in this channel group. The possible values for a channel group mode are as follows.

- ON (default)—The member ports only operate as part of a PortChannel or remain inactive. In this mode, the PortChannel protocol is not initiated. However, if a PortChannel protocol frame is received from a peer port, the software indicates its nonnegotiable status. This mode is backward compatible with the existing implementation of PortChannels in releases prior to Release 2.0(1b), where the channel group mode is implicitly assumed to be ON. In Cisco MDS SAN-OS Releases 1.3 and earlier, the only available PortChannel mode was the ON mode. PortChannels configured in the ON Mode require you to explicitly enable and disable the PortChannel member ports at either end if you add or remove ports from the PortChannel configuration. You must physically verify that the local and remote ports are connected to each other.
- ACTIVE—The member ports initiate PortChannel protocol negotiation with the peer port(s) regardless of the channel group mode of the peer port. If the peer port, while configured in a channel group, does not support the PortChannel protocol, or responds with a nonnegotiable status, it will default to the ON mode behavior. The ACTIVE PortChannel mode allows automatic recovery without explicitly enabling and disabling the PortChannel member ports at either end.

Table 21-1 compares ON and ACTIVE modes.

Table 21-1 Channel Group Configuration Differences (contd.)

ON Mode	ACTIVE Mode
No protocol is exchanged.	A PortChannel protocol negotiation is performed with the peer ports.
Moves interfaces to the suspended state if its operational values are incompatible with the PortChannel.	Moves interfaces to the isolated state if its operational values are incompatible with the PortChannel.

Table 21-1 Channel Group Configuration Differences (contd.)

ON Mode	ACTIVE Mode
When you add or modify a PortChannel member port configuration, you must explicitly disable (shut) and enable (no shut) the PortChannel member ports at either end.	When you add or modify a PortChannel interface, the PortChannel automatically recovers.
Port initialization is not synchronized.	There is synchronized startup of all ports in a channel across peer switches.
All misconfigurations are not detected as no protocol is exchanged.	Consistently detect misconfigurations using a PortChannel protocol.
Transitions misconfigured ports to the suspended state. You must explicitly disable (shut) and enable (no shut) the member ports at either end.	Transitions misconfigured ports to the isolated state to correct the misconfiguration. Once you correct the misconfiguration, the protocol ensures automatic recovery.
This is the default mode.	You must explicitly configure this mode.

Configuring Port Channel Modes

To configure active mode using Fabric Manager, follow these steps:

Step 1 Expand **ISLs** and then select **Port Channels** in the Physical Attributes pane.

You see the PortChannels configured in the Information pane.

Step 2 Click the **Protocols** tab and, from the Mode drop-down menu, select the appropriate mode for the Port Channel.

Step 3 Click **Apply Changes** icon to save any modifications or click **Undo Changes** to discard any changes.

About PortChannel Deletion

When you delete the PortChannel, the corresponding channel membership is also deleted. All interfaces in the deleted PortChannel convert to individual physical links. After the PortChannel is removed, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down.

If you delete the PortChannel for one port, then the individual ports within the deleted PortChannel retain the compatibility parameter settings (speed, mode, port VSAN, allowed VSAN, and port security). You can explicitly change those settings as required.

- If you use the default ON mode, to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, then the PortChannel ports automatically recover from the deletion.

Deleting PortChannels

To delete a PortChannel using the PortChannel Wizard in Fabric Manager, follow these steps:

Step 1 Click the **PortChannel Wizard** icon in the toolbar (see Figure 21-13).

Figure 21-13 PortChannel Wizard Icon

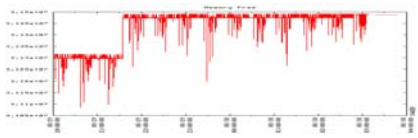

You see the first PortChannel Wizard screen.

Step 2 Select the existing PortChannel that you want to delete and click **Next**. You see a list of the ISLs currently associated with this PortChannel.

Step 3 Click **Next**. You see an editable list of associated ISLs and available ISLs for this PortChannel.

Step 4 Click each associated ISL and click the **left arrow** to remove all ISLs from the PortChannel.

Step 5 Check the **Delete Port Channel If Empty** check box to delete this PortChannel.

Step 6 Click **Finish** to save any modifications or click **Cancel** to discard any changes.

Interfaces in a PortChannel

You can add or remove a physical interface (or a range of interfaces) to an existing PortChannel. The compatible parameters on the configuration are mapped to the PortChannel. Adding an interface to a PortChannel increases the channel size and bandwidth of the PortChannel. Removing an interface from a PortChannel decreases the channel size and bandwidth of the PortChannel.

This section describes interface configuration for a PortChannel and includes the following topics:

About Interface Addition to a PortChannel

You can add a physical interface (or a range of interfaces) to an existing PortChannel. The compatible parameters on the configuration are mapped to the PortChannel. Adding an interface to a PortChannel increases the channel size and bandwidth of the PortChannel.

After the members are added, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the "32-Port Switching Module Configuration Guidelines" section on page 21-3).

Compatibility Check

A compatibility check ensures that the same parameter settings are used in all physical ports in the channel. Otherwise, they cannot become part of a PortChannel. The compatibility check is performed before a port is added to the PortChannel.

The check ensures that the following parameters and settings match at both ends of a PortChannel:

- Capability parameters (type of interface, Gigabit Ethernet at both ends, or Fibre Channel at both ends).
- Administrative compatibility parameters (speed, mode, port VSAN, allowed VSAN, and port security).
- Operational parameters (speed and remote switch's WWN).

A port addition procedure fails if the capability and administrative parameters in the remote switch are incompatible with the capability and administrative parameters in the local switch. If the compatibility check is successful, the interfaces are operational and the corresponding compatibility parameter settings apply to these interfaces.

Suspended and Isolated States

If the operational parameters are incompatible, the compatibility check fails and the interface is placed in a suspended or isolated state based on the configured mode:

- An interface enters the suspended state if the interface is configured in the ON mode.
- An interface enters the isolated state if the interface is configured in the ACTIVE mode.

Adding an Interface to a PortChannel

To add an interface or range of interfaces to a PortChannel using Fabric Manager, follow these steps:

Step 1 Expand **ISLs** and then select **Port Channels** in the Physical Attributes pane.

You see the PortChannels configured in the Information pane (see Figure 21-14).

Figure 21-14 Port Channels

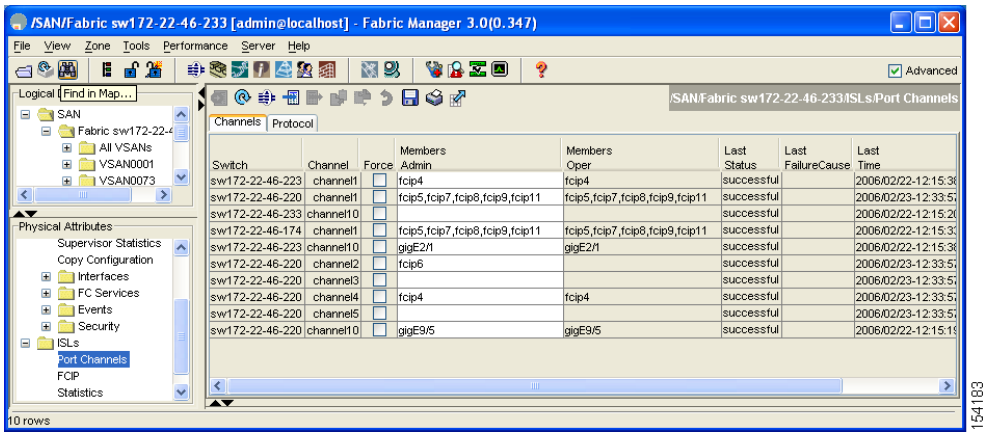

Step 2 Click the **Channels** tab and find the switch and PortChannel that you want to edit.

Step 3 Set Members Admin to the interface or list of interfaces that you want to add to the PortChannel.

Step 4 Click the **Apply Changes** icon to save any modifications or click **Undo Changes** to discard any changes.

About Forcing an Interface Addition

You can force the port configuration to be overwritten by the PortChannel. In this case, the interface is added to a PortChannel.

- If you use the default ON mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, then the PortChannel ports automatically recover from the addition.

Note When PortChannels are created from within an interface, the **force** option cannot be used.

After the members are forcefully added, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the "32-Port Switching Module Configuration Guidelines" section on page 21-3).

Forcing an Interface Addition

To force the addition of a port to a PortChannel using Fabric Manager, follow these steps:

Step 1 Expand **ISLs** and then select **Port Channels** in the Physical Attributes pane. You see the PortChannels configured in the Information pane.

Step 2 Click the **Channels** tab and find the switch and PortChannel that you want to edit.

Step 3 Set Members Admin to the interface or list of interfaces that you want to add to the PortChannel.

Step 4 Check the **Force** check box to force this interface addition.

Step 5 Click the **Apply Changes** to save any modifications or click **Undo Changes** to discard any changes.

About Interface Deletion from a PortChannel

When a physical interface is deleted from the PortChannel, the channel membership is automatically updated. If the deleted interface is the last operational interface, then the PortChannel status is changed to a down state. Deleting an interface from a PortChannel decreases the channel size and bandwidth of the PortChannel.

- If you use the default ON mode, to avoid inconsistent states across switches, and to maintain consistency across switches, the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, the PortChannel ports automatically recover from the deletion.

After the members are deleted, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the "32-Port Switching Module Configuration Guidelines" section on page 21-3).

Deleting Interfaces from a PortChannel

To delete a physical interface (or a range of physical interfaces) using Fabric Manager, follow these steps:

Step 1 Expand **ISLs** and then select **Port Channels** in the Physical Attributes pane.

You see the PortChannels configured in the Information pane.

Step 2 Click the **Channels** tab and find the switch and PortChannel that you want to edit.

Step 3 Remove the interface or list of interfaces you want deleted in the Members the Admin column.

Step 4 Click **Apply Changes** to save any modifications or click **Undo Changes** to discard any changes.

PortChannel Protocol

In earlier Cisco SAN-OS releases, PortChannels required additional administrative tasks to support synchronization. The Cisco SAN-OS software provides robust error detection and synchronization capabilities. You can manually configure channel groups or they can be automatically created. In both cases, the channel groups have the same capability and configurational parameters. Any change in configuration applied to the associated PortChannel interface is propagated to all members of the channel group.

A protocol to exchange PortChannel configurations is available in all Cisco MDS switches. This addition simplifies PortChannel management with incompatible ISLs. An additional autocreation mode enables ISLs with compatible parameters to automatically form channel groups without manual intervention.

The PortChannel protocol is enabled by default.

The PortChannel protocol expands the PortChannel functional model in Cisco MDS switches. It uses the exchange peer parameters (EPP) services to communicate across peer ports in an ISL. Each switch uses the information received from the peer ports along with its local configuration and operational values to decide if it should be part of a PortChannel. The protocol ensures that a set of ports are eligible to be part of the same PortChannel. They are only eligible to be part of the same port channel if all the ports have a compatible partner.

The PortChannel protocol uses two sub-protocols:

- Bringup protocol—Automatically detects misconfigurations so you can correct them. This protocol synchronizes the PortChannel at both ends so that all frames for a given flow (as identified by the source FC ID, destination FC ID and OX_ID) are carried over the same physical link in both directions. This helps make applications like write acceleration work for PortChannels over FCIP links.
- Autocreation protocol—Automatically aggregates compatible ports into a PortChannel.

This section describes how to configure the PortChannel protocol and includes the following sections:

About Channel Group Creation

Assuming link A1-B1 comes up first in Figure 21-15, that link is operational as an individual link.When the next link, say A2-B2 comes up, the PortChannel protocol identifies if this link is compatible with link A1-B1 and automatically creates channel groups 10 and 20 in the respective switches. If link A3-B3 can join the channel groups (and hence, the PortChannels), the respective ports have compatible configurations. If link A4-B4 operates as an individual link, it is because of the incompatible configuration of the two end ports with the other member ports in this channel group.

Figure 21-15 Autocreating Channel Groups

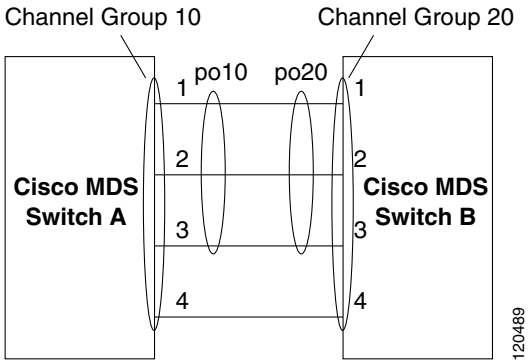

The channel group numbers are selected dynamically, and as such, the administrative configuration of the ports forming the channel group at either end are applicable to the newly created channel group. The channel group number being chosen dynamically may be different across reboots for the same set of PortChannels based on the order of ports that are initialized in the switch.

Table 21-2 identifies the differences between user-configured and auto-configured channel groups.

Table 21-2 Channel Group Configuration Differences

User-Configured Channel Group	Autocreated Channel Group
Manually configured by the user.	Created automatically when compatible links come up between two compatible switches, if channel group autocreation is enabled in all ports at both ends.
Member ports cannot participate in autocreation of channel groups. The autocreation feature cannot be configured.	None of these ports are members of a user-configured channel group.
You can form the PortChannel with a subset of the ports in the channel group. Incompatible ports remain in a suspended or isolated state depending on the ON or ACTIVE mode configuration.	All ports included in the channel group participate in the PortChannel—no member port becomes isolated or suspended; instead, the member port is removed from the channel group when the link is found to be incompatible.

Table 21-2 Channel Group Configuration Differences (continued)

User-Configured Channel Group	Autocreated Channel Group
Any administrative configuration made to the PortChannel is applied to all ports in the channel group, and you can save the configuration for the PortChannel interface.	Any administrative configuration made to the PortChannel is applied to all ports in the channel group, but the configurations are saved for the member ports; no configuration is saved for the PortChannel interface. You can explicitly convert this channel group, if required.
You can remove any channel group and add members to a channel group.	You cannot remove a channel group, or add/remove any of its members. The channel group is removed when no member ports exist.

About Autocreation

The autocreation protocol has the following functionality:

- A port is not allowed to be configured as part of a PortChannel when the autocreation feature is enabled. These two configurations are mutually exclusive.
- Autocreation must be enabled in both the local and peer ports to negotiate a PortChannel.
- Aggregation occurs in one of two ways:
  - A port is aggregated into a compatible autocreated PortChannel.
  - A port is aggregated with another compatible port to form a new PortChannel.
- Newly created PortChannels are allocated from the maximum possible PortChannel (128) in a decreasing order based on availability. If all 128 numbers are used up, aggregation is not allowed.
- You cannot change the membership or delete an autocreated PortChannel.
- When you disable autocreation, all member ports are removed from the autocreated PortChannel.
- Once the last member is removed from an autocreated PortChannel, the channel is automatically deleted and the number is released for reuse.
- An autocreated PortChannel is not persistent through a reboot. An autocreated PortChannel can be manually configured to appear the same as a persistent PortChannel. Once the PortChannel is made persistent, the autocreation feature is disabled in all member ports.

Tip When enabling autocreation in any switch in the Cisco MDS 9000 Family, we recommend that you retain at least one interconnected port between the switches without any autocreation configuration. If all ports between two switches are configured with the autocreation feature at the same time, you may face a possible traffic disruption between these two switches as the ports are automatically disabled and reenabled when ports are added to an autocreated PortChannel.

Enabling and Configuring Autocreation

To configure PortChannel autocreation, check the **Dynamically form Port Channel Group from selected ISLs** option in the PortChannel Wizard. See Configuring PortChannels, page 21-9.

About Manually Configured Channel Groups

A user-configured channel group cannot be converted to an autocreated channel group. However, you can convert an autocreated channel group to a manual channel group. Once performed, this task is irreversible— the channel group number does not change, but the member ports operate according to the properties of the manually configured channel group, and the autocreation of channel group is implicitly disabled for all member ports.

Tip If you enable persistence, be sure to enable it at both ends of the PortChannel.

Converting to Manually Configured Channel Groups

To convert an autocreated channel group to a user-configured channel group using Fabric Manager, follow these steps:

Step 1 Expand **ISLs** and then select **Port Channels** in the Physical Attributes pane.

You see the PortChannels configured in the Information pane.

Step 2 Click the **Protocol** tab.

You see the switch protocols shown in Figure 21-16.

Figure 21-16 Switch Protocols

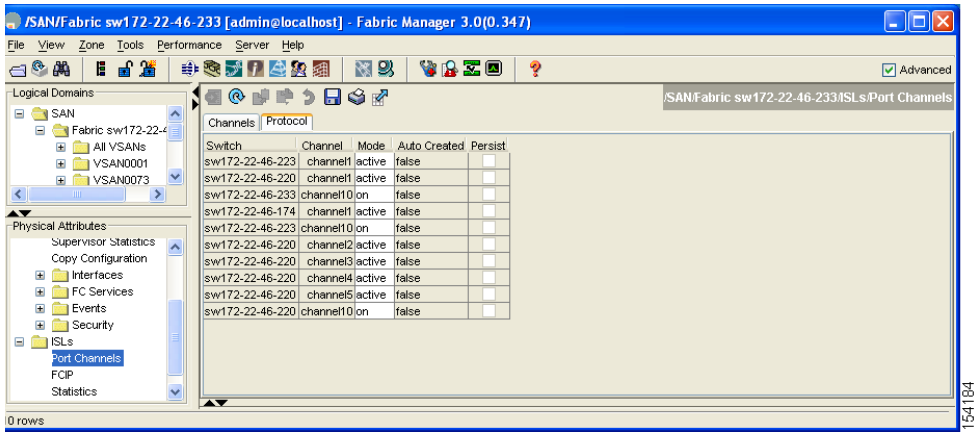

Step 3 Check the **Persist** check box for each channel that you want to convert to a manually configured channel group.

Step 4 Click **Apply Changes** to save any modifications or click **Undo Changes** to discard any changes.

PortChannel Configuration Verification

You can use the Information pane in Fabric Manager to verify your PortChannel Configuration. See Figure 21-17.

Figure 21-17 PortChannel Summary in Fabric Manager

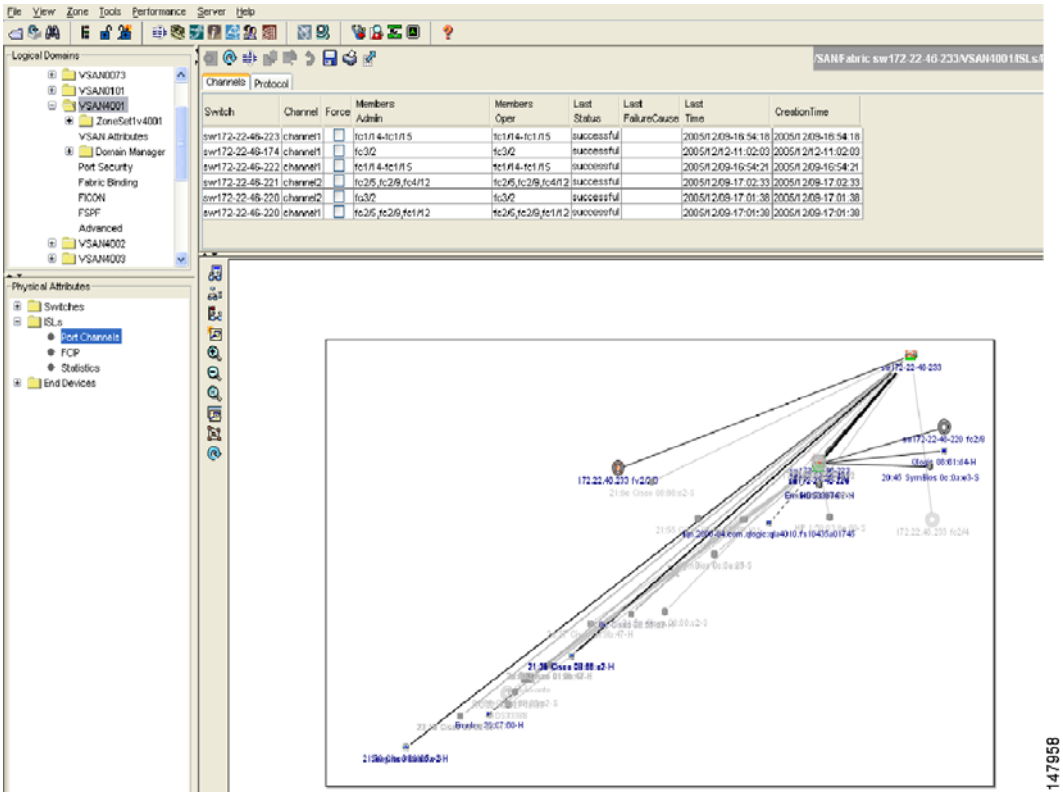

Default Settings

Table 21-3 lists the default settings for PortChannels.

Table 21-3 Default PortChannel Parameters

Parameters	Default
PortChannels	FSPF is enabled by default.
Create PortChannel	Administratively up.
Default PortChannel mode	On.
Autocreation	Disabled.

CHAPTER 22

Configuring Domain Parameters

The Fibre Channel domain (fcdomain) feature performs principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions as described in the FC-SW-2 standards. The domains are configured on a per VSAN basis. If you do not configure a domain ID, the local switch uses a random ID.

Caution Changes to fcdomain parameters should not be performed on a daily basis. These changes should be made by an administrator or individual who is completely familiar with switch operations.

Tip When you change the configuration, be sure to save the running configuration. The next time you reboot the switch, the saved configuration is used. If you do not save the configuration, the previously saved startup configuration is used.

This chapter includes the following sections:

Fibre Channel Domains

This section describes each fcdomain phase:

- Principal switch selection—This phase guarantees the selection of a unique principal switch across the fabric.
- Domain ID distribution—This phase guarantees each switch in the fabric obtains a unique domain ID.
- FC ID allocation—This phase guarantees a unique FC ID assignment to each device attached to the corresponding switch in the fabric.
- Fabric reconfiguration—This phase guarantees a resynchronization of all switches in the fabric to ensure they simultaneously restart a new principal switch selection phase.

See Figure 22-1.

Figure 22-1 Sample fcdomain Configuration

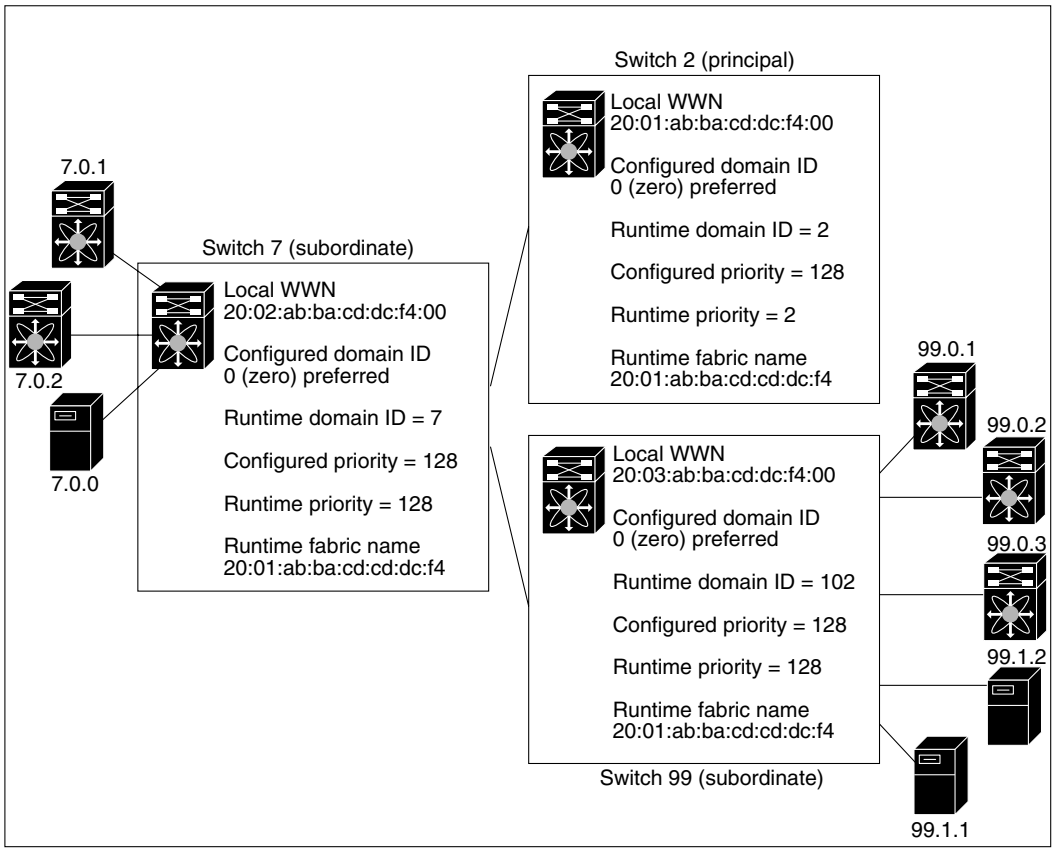

Note Domain IDs and VSAN values used in all procedures are only provided as examples. Be sure to use IDs and values that apply to your configuration.

This section describes the fcdomain feature and includes the following topics:

About Domain Restart

Fibre Channel domains can be started disruptively or nondisruptively. If you perform a disruptive restart, reconfigure fabric (RCF) frames are sent to other switches in the fabric and data traffic is disrupted on all the switches in the VSAN. If you perform a nondisruptive restart, build fabric (BF) frames are sent to other switches in the fabric and data traffic is disrupted only on the switch.

Note A static domain is specifically configured by the user and may be different from the runtime domain. If the domain IDs are different, the runtime domain ID changes to take on the static domain ID after the next restart, either disruptive or non-disruptive.

Tip If a VSAN is in interop mode, you cannot restart the fcdomain for that VSAN disruptively.

You can apply most of the configurations to their corresponding runtime values. Each of the following sections provide further details on how the fcdomain parameters are applied to the runtime values.

Restarting a Domain

To restart the fabric disruptively or nondisruptively using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx*** and then select **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to restart.

You see the Running tab configuration of the domain in the Information pane.

Figure 22-2 Running Domain Configuration

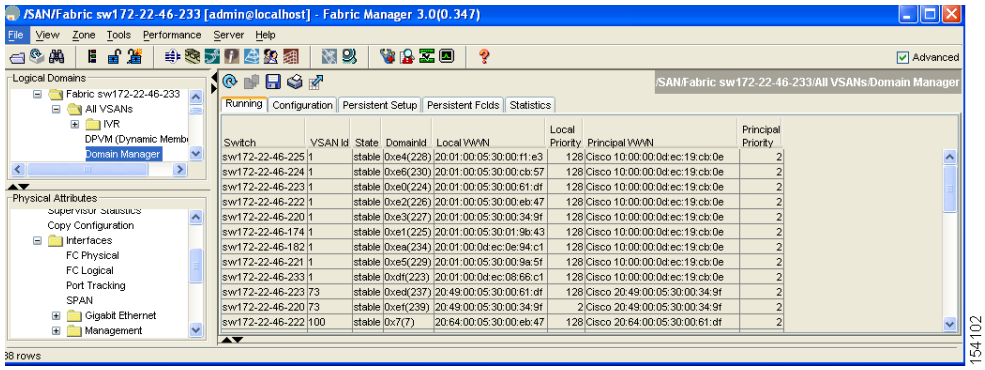

Step 2 Click the **Configuration** tab.

You see the switch configuration shown in Figure 22-3.

Figure 22-3 Configuring Domains

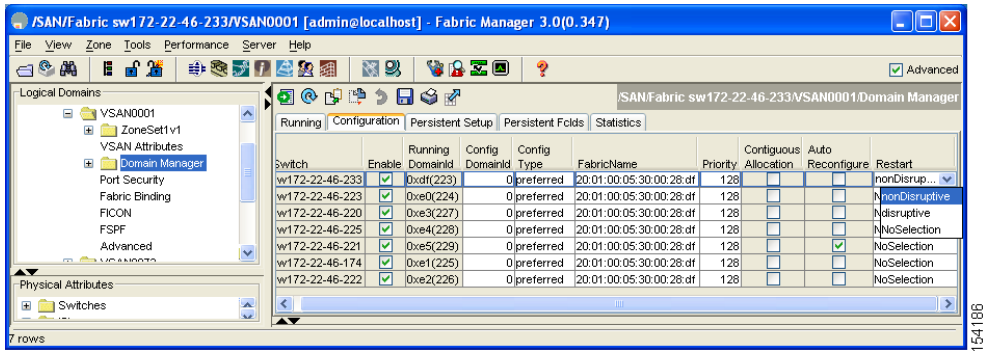

Step 3 Set the Restart drop-down menu to **disruptive** or **nonDisruptive** for any switch in the fabric that you want to restart the fcdomain.

Step 4 Click **Apply Changes** to issue this fcdomain restart or click **Undo Changes** to drop any unsaved changes.

About Switch Priority

By default, the configured priority is 128. The valid range to set the priority is between 1 and 254. Priority 1 has the highest priority. Value 255 is accepted from other switches, but cannot be locally configured.

Any new switch cannot become the principal switch when it joins a stable fabric. During the principal switch selection phase, the switch with the highest priority becomes the principal switch. If two switches have the same configured priority, the switch with the lower WWN becomes the principal switch.

The priority configuration is applied to runtime when the fcdomain is restarted (see the "About Domain Restart" section on page 22-3). This configuration is applicable to both disruptive and nondisruptive restarts.

Configuring Switch Priority

To configure the priority for the principal switch using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx*** and then select **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to set the principal switch priority for.

You see the domain's running configuration in the Information pane shown in Figure 22-4.

Figure 22-4 Running Domain Configuration

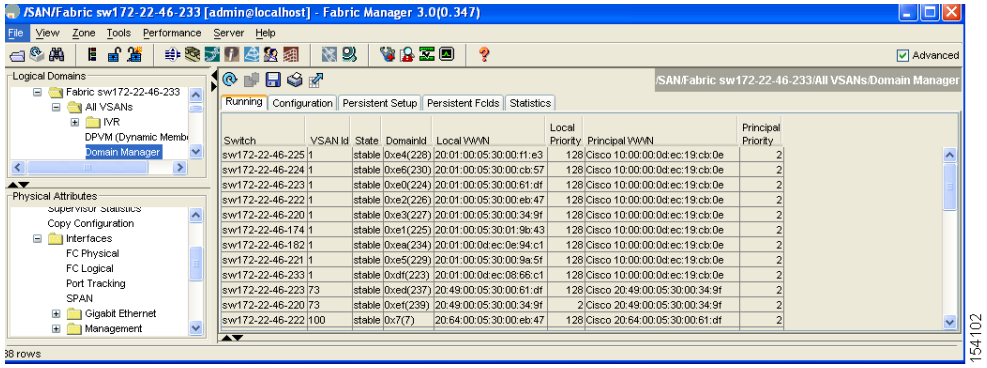

Step 2 Set Priority to a high value for the switch in the fabric that you want to be the principal switch.

Step 3 Click **Apply Changes** to save these changes or click **Undo Changes** to drop any unsaved changes.

About fcdomain Initiation

By default, the fcdomain feature is enabled on each switch. If you disable the fcdomain feature in a switch, that switch can no longer participate with other switches in the fabric. The fcdomain configuration is applied to runtime through a disruptive restart.

Enabling or Disabling fcdomains

To disable fcdomains in a single VSAN or a range of VSANs using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx*** and then select **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to disable fcdomain for.

You see the domain's running configuration in the Information pane.

Step 2 Click the **Configuration** tab and uncheck the **Enable** check box shown in Figure 22-5 for each switch in the fabric that you want to disable fcdomain on.

Figure 22-5 Configuring Domains

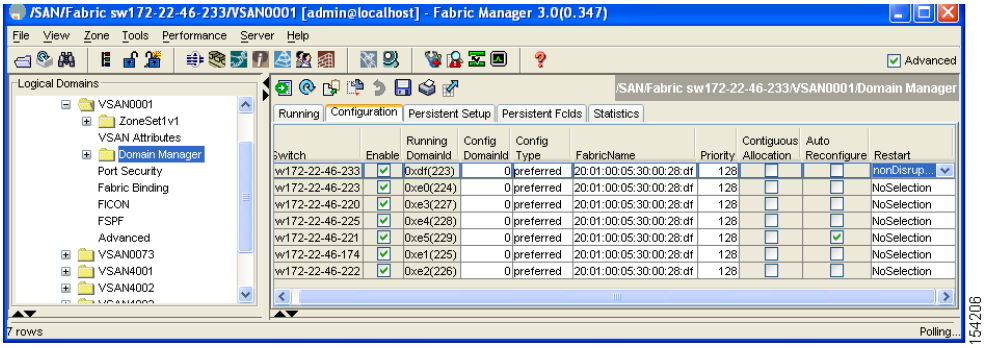

Step 3 Click **Apply Changes** to save these changes or click **Undo Changes** to drop any unsaved changes.

About Fabric Names

By default the configured fabric name is 20:01:00:05:30:00:28:df.

- When the fcdomain feature is disabled, the runtime fabric name is the same as the configured fabric name.
- When the fcdomain feature is enabled, the runtime fabric name is the same as the principal switch's WWN.

The fabric name is applied to runtime through a disruptive restart when the fcdomain is configured as disabled (see the "About Domain Restart" section on page 22-3).

Setting Fabric Names

To set the fabric name value for a disabled fcdomain using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx*** and then select **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to set the fabric name for.

You see the running configuration of the domain in the Information pane.

Step 2 Click the **Configuration** tab and set the fabric name for each switch in the fabric.

Step 3 Click **Apply Changes** to save these changes or click **Undo Changes** to drop any unsaved changes.

About Incoming RCFs

You can choose to reject RCF request frames on a per-interface, per-VSAN basis. By default, the RCF reject option is disabled (that is, RCF request frames are not automatically rejected).

The RCF reject option takes immediate effect to runtime through a disruptive restart (see the "About Domain Restart" section on page 22-3).

Rejecting Incoming RCFs

To reject incoming RCF request frames using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** and then select **FC Physical** in the Physical Attributes pane**.**

You see the Fibre Channel configuration in the Information pane.

Step 2 Click the **Domain Mgr** tab.

You see the information in Figure 22-6.

Figure 22-6 Rejecting Incoming RCF Request Frames

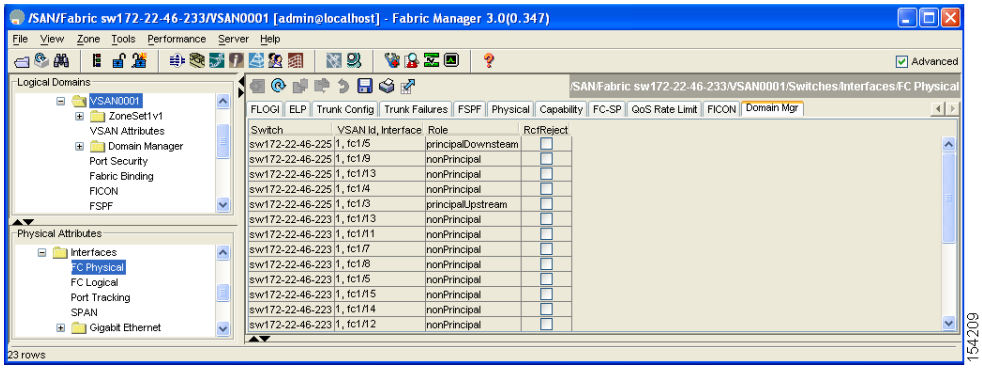

Step 3 Check the **RcfReject** check box for each interface that you want to reject RCF request frames on.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to drop any unsaved changes.

About Autoreconfiguring Merged Fabrics

By default, the autoreconfigure option is disabled. When you join two switches belonging to two different stable fabrics that have overlapping domains, the following cases apply:

- If the autoreconfigure option is enabled on both switches, a disruptive reconfiguration phase is started.
- If the autoreconfigure option is disabled on either or both switches, the links between the two switches become isolated.

The autoreconfigure option takes immediate effect at runtime. You do not need to restart the fcdomain. If a domain is currently isolated due to domain overlap, and you later enable the auto-reconfigure option on both switches, the fabric continues to be isolated. If you enabled the auto-reconfigure option on both switches before connecting the fabric, a disruptive reconfiguration (RCF) will occur. A disruptive reconfiguration may affect data traffic. You can nondisruptively reconfigure the fcdomain by changing the configured domains on the overlapping links and getting rid of the domain overlap.

Enabling Autoreconfiguration

To enable automatic reconfiguration in a specific VSAN (or range of VSANs) using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx*** and then select **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to enable automatic reconfiguration for.

You see the running configuration of the domain in the Information pane.

Step 2 Select the **Configuration** tab and check the **Auto Reconfigure** check box for each switch in the fabric that you want to automatically reconfigure.

Step 3 Click **Apply Changes** to save these changes or click **Undo Changes** to drop any unsaved changes.

Domain IDs

Domain IDs uniquely identify a switch in a VSAN. A switch may have different domain IDs in different VSANs. The domain ID is part of the overall FC ID.

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

About Domain IDs

The configured domain ID can be preferred or static. By default, the configured domain ID is 0 (zero) and the configured type is preferred. If you do not configure a domain ID, the local switch sends a random ID in its request.

When a subordinate switch requests a domain, the following process takes place (see Figure 22-7):

1. The local switch sends a configured domain ID request to the principal switch.
2. The principal switch assigns the requested domain ID if available. Otherwise, it assigns another available domain ID.

Figure 22-7 Configuration Process Using the preferred Option

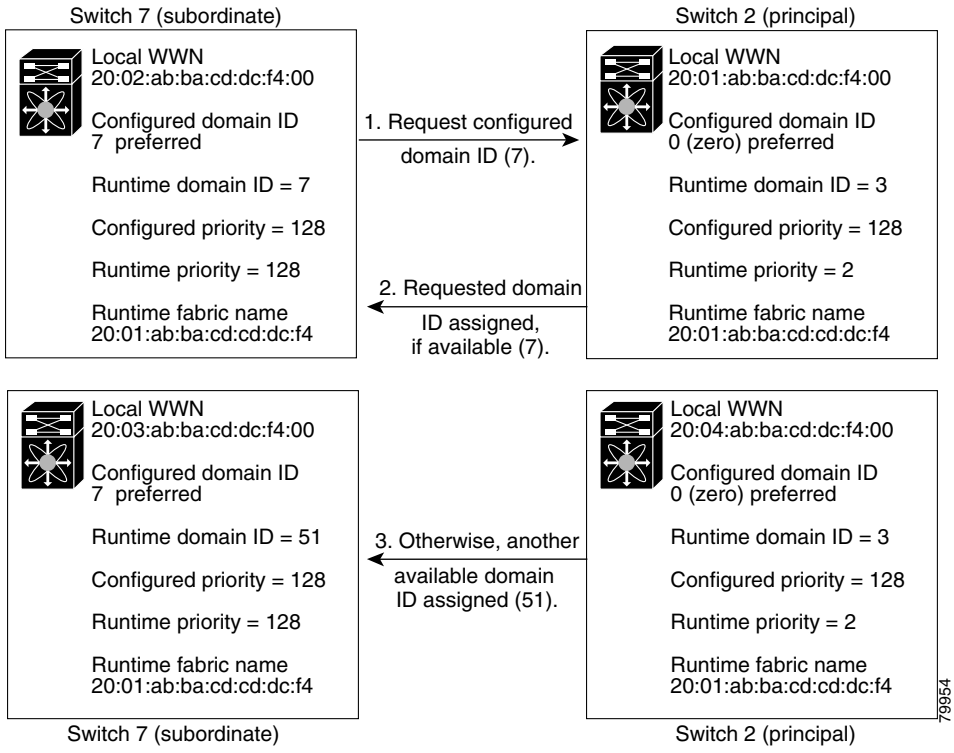

The behavior for a subordinate switch changes based on three factors:

- The allowed domain ID lists.
- The configured domain ID.
- The domain ID that the principal switch has assigned to the requesting switch.

In specific situations, the changes are as follows:

- When the received domain ID is not within the allowed list, the requested domain ID becomes the runtime domain ID and all interfaces on that VSAN are isolated.
- When the assigned and requested domain IDs are the same, the preferred and static options are not relevant, and the assigned domain ID becomes the runtime domain ID.
- When the assigned and requested domain IDs are different, the following cases apply:
  - If the configured type is static, the assigned domain ID is discarded, all local interfaces are isolated, and the local switch assigns itself the configured domain ID, which becomes the runtime domain ID.
  - If the configured type is preferred, the local switch accepts the domain ID assigned by the principal switch and the assigned domain ID becomes the runtime domain ID.

If you change the configured domain ID, the change is only accepted if the new domain ID is included in all the allowed domain ID lists currently configured in the VSAN. Alternatively, you can also configure zero-preferred domain ID.

Note The 0 (zero) value can be configured only if you use the preferred option.

While the static option can be applied at runtime after a disruptive or nondisruptive restart, the preferred option is applied at runtime only after a disruptive restart (see the "About Domain Restart" section on page 22-3).

Tip When the FICON feature is enabled in a given VSAN, the domain ID for that VSAN remains in the static state. You can change the static ID value but you cannot change it to the preferred option.

Note In an IVR without NAT configuration, if one VSAN in the IVR topology is configured with static domain IDs, then the other VSANs (edge or transit) in the topology should has to be configured with static domain IDs.

Note In an IVR NAT configuration, if one VSAN in the IVR topology is configured with static domain IDs, then the IVR domains that can be exported to that VSAN must also be assigned static domains.

Caution You must restart the fcdomain if you want to apply the configured domain changes to the runtime domain.

Specifying Static or Preferred Domain IDs

To specify a preferred or a static domain ID using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx*** and then select **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to configure the domain ID for.

You see the running configuration of the domain in the Information pane.

Step 2 Enter a value for the Config DomainID and click **static** or **preferred** from the Config Type drop-down menu to set the domain ID for switches in the fabric.

Step 3 Click **Apply Changes** to save these changes or click **Undo Changes** to drop any unsaved changes.

About Allowed Domain ID Lists

By default, the valid range for an assigned domain ID list is from 1 to 239. You can specify a list of ranges to be in the allowed domain ID list and separate each range with a comma. The principal switch assigns domain IDs that are available in the locally configured allowed domain list.

Use allowed domain ID lists to design your VSANs with non-overlapping domain IDs. This helps you in the future if you need to implement IVR without the NAT feature.

Tip If you configure an allowed list on one switch in the fabric, we recommend you configure the same list in all other switches in the fabric to ensure consistency.

An allowed domain ID list must satisfy the following conditions:

- If this switch is a principal switch, all the currently assigned domain IDs must be in the allowed list.
- If this switch is a subordinate switch, the local runtime domain ID must be in the allowed list.
- The locally configured domain ID of the switch must be in the allowed list.
- The intersection of the assigned domain IDs with other already configured domain ID lists must not be empty.

Configuring Allowed Domain ID Lists

To configure the allowed domain ID list using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx* > Domain Manager** and then select **Allowed** in the Logical Domains pane for the fabric and VSAN that you want to set the allowed domain ID list for.

You see the CFS configuration in the Information pane. See Figure 22-8.

Figure 22-8 Allowed CFS Configuration Information

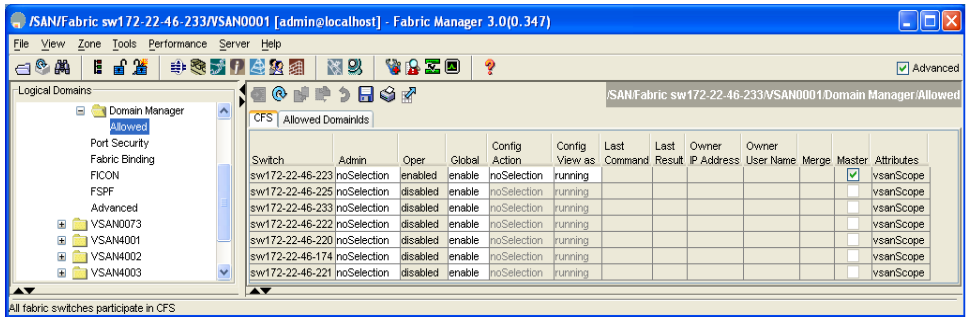

Step 2 Set the Admin drop-down menu to **enable** and set the Global drop-down menu to **enable.**

Step 3 Click **Apply Changes** to enable CFS distribution for the allowed domain ID list.

Step 4 Select the **Allowed DomainIds** tab.

You see the Allowed Domain ID screen shown in Figure 22-9.

Figure 22-9 Allowed Domain ID List

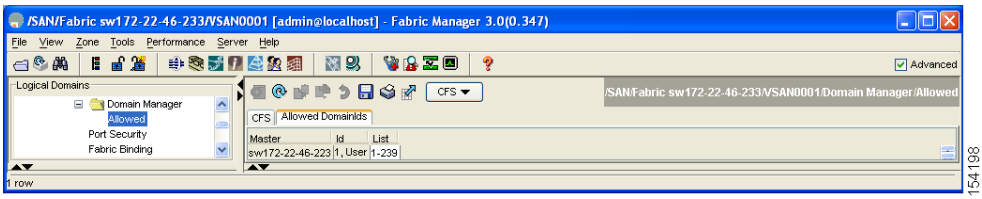

Step 5 Set the list to the allowed domain IDs list for this domain.

Step 6 Select the **CFS** tab and set Config Action to **commit**.

Step 7 Click **Apply Changes** to commit this allowed domain ID list and distribute it throughout the VSAN or click **Undo Changes** to drop any unsaved changes.

About CFS Distribution of Allowed Domain ID Lists

As of Cisco SAN-OS Release 3.0(1), you can enable the distribution of the allowed domain ID list s configuration information to all Cisco MDS switches in the fabric using the Cisco Fabric Services (CFS) infrastructure. This feature allows you to synchronize the configuration across the fabric from the console of a single MDS switch. Since the same configuration is distributed to the entire VSAN, you avoid possible misconfiguration and the likelihood that two switches in the same VSAN have configured incompatible allowed domains.

Note All switches in the fabric must be running Cisco SAN-OS Release 3.0(1) or later to distribute the allowed domain ID list using CFS.

Use CFS to distribute the allowed domain ID list to ensure consistency in the allowed domain ID lists on all switches in the VSAN.

For more information about CFS, see Chapter 12, “Using the CFS Infrastructure.”

Enabling Distribution

CFS distribution of allowed domain ID lists is disabled by default.

To enable (or disable) allowed domain ID list configuration distribution using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx* > Domain Manager** and then select **Allowed** in the Logical Domains pane for the fabric and VSAN that you want to set the allowed domain ID list for.

You see the CFS configuration in the Information pane.

Step 2 Set the Admin drop-down menu to **enable** and the Global drop-down menu to **enable** to enable CFS distribution for the allowed domain ID list.

Step 3 Click **Apply Changes** to enable CFS distribution for the allowed domain ID list or click **Undo Changes** to drop any unsaved changes.

Locking the Fabric

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

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

Committing Changes

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

To commit pending domain configuration changes and release the lock using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx* > Domain Manager** and then select **Allowed** in the Logical Domains pane for the fabric and VSAN that you want to set the allowed domain ID list for.

You see the CFS configuration in the Information pane.

Step 2 Set the Config Action drop-down menu to **commit**.

Step 3 Click **Apply Changes** to commit the allowed domain ID list and distribute it throughout the VSAN or click **Undo Changes** to drop any unsaved changes.

Discarding Changes

At any time, you can discard the pending changes to the domain configuration and release the fabric lock. If you discard (abort) the pending changes, the configuration remains unaffected and the lock is released.

To discard pending domain configuration changes and release the lock using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx* > Domain Manager** and then select **Allowed** in the Logical Domains pane for the fabric and VSAN that you want to set the allowed domain ID list for.

You see the CFS configuration in the Information pane.

Step 2 Set the Config Action drop-down menu to **abort**.

Step 3 Click **Apply Changes** to discard any pending changes to the allowed domain ID list or click **Undo Changes** to drop any unsaved changes.

Clearing a Fabric Lock

If you have performed a domain configuration task and have not released the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your pending changes are discarded and the fabric lock is released.

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

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

Step 1 Expand **Fabric*xx* > VSAN*xx* > Domain Manager** and then select **AllowedId** in the Logical Domains pane for the fabric and VSAN for which you want the allowed domain ID list.

You see the CFS configuration in the Information pane.

Step 2 Set the Config Action drop-down menu to **clear**.

Step 3 Click **Apply Changes** to clear the fabric lock or click **Undo Changes** to drop any unsaved changes.

Displaying Pending Changes

To display the pending configuration changes using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx* > Domain Manager > Allowed** in the Logical Domains pane for the fabric and VSAN that you want to set the allowed domain ID list for.

You see the CFS configuration in the Information pane.

Step 2 Set the Config View As drop-down menu to **pending**.

Step 3 Click **Apply Changes** to clear the fabric lock or click **Undo Changes** to drop any unsaved changes.

Step 4 Select the **AllowedDomainIds** tab.

You see the pending configuration for the allowed domain IDs list.

Displaying Session Status

To display the status of the distribution session using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx* > Domain Manager** and then select **Allowed** in the Logical Domains pane for the fabric and VSAN for which you want to set the allowed domain ID list.

Step 2 View the CFS configuration and session status in the Information pane.

About Contiguous Domain ID Assignments

By default, the contiguous domain assignment is disabled. When a subordinate switch requests the principal switch for two or more domains and the domains are not contiguous, the following cases apply:

- If the contiguous domain assignment is enabled in the principal switch, the principal switch locates contiguous domains and assigns them to the subordinate switches. If contiguous domains are not available, the SAN-OS software rejects this request.
- If the contiguous domain assignment is disabled in the principal switch, the principal switch assigns the available domains to the subordinate switch.

Enabling Contiguous Domain ID Assignments

To enable contiguous domains in a specific VSAN (or a range of VSANs) using Fabric Manager, follow these steps

Step 1 Expand **Fabric*xx* > VSAN*xx*** and then select **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to enable contiguous domains for.

You see the running configuration of the domain in the Information pane.

Step 2 Select the **Configuration** tab and check the **Contiguous Allocation** check box for each switch in the fabric that will have contiguous allocation.

Step 3 Click **Apply Changes** to save these changes or click **Undo Changes** to drop any unsaved changes.

FC IDs

When an N or NL port logs into a Cisco MDS 9000 Family switch, it is assigned an FC ID. By default, the persistent FC ID feature is enabled. If this feature is disabled, the following consequences apply:

- An N or NL port logs into a Cisco MDS 9000 Family switch. The WWN of the requesting N or NL port and the assigned FC ID are retained and stored in a volatile cache. The contents of this volatile cache are not saved across reboots.
- The switch is designed to preserve the binding FC ID to the WWN on a best-effort basis. For example, if one N port disconnects from the switch and its FC ID is requested by another device, this request is granted and the WWN with the initial FC ID association is released.

- The volatile cache stores up to 4000 entries of WWN to FC ID binding. If this cache is full, a new (more recent) entry overwrites the oldest entry in the cache. In this case, the corresponding WWN to FC ID association for the oldest entry is lost.
- The switch connection behavior differs between N ports and NL ports:
  - N ports receive the same FC IDs if disconnected and reconnected to any port within the same switch (as long as it belongs to the same VSAN).
  - NL ports receive the same FC IDs only if connected back to the same port on the switch to which they were originally connected.

This section describes configuring FC IDs and includes the following topics:

About Persistent FC IDs

If this feature remains enabled, the following consequences apply:

- The currently *in use* FC IDs in the fcdomain are saved across reboots.
- The fcdomain automatically populates the database with dynamic entries that the switch has learned about after a device (host or disk) is plugged into a port interface.

Note If you connect to the switch from an AIX or HP-UX host, be sure to enable the persistent FC ID feature in the VSAN that connects these hosts.

Note FC IDs are enabled by default. This change of default behavior from releases prior to Cisco MDS SAN-OS Release 2.0(1b) prevents FC IDs from being changed after a reboot. You can disable this option for each VSAN.

A persistent FC ID assigned to an F port can be moved across interfaces and can continue to maintain the same persistent FC ID.

Note Persistent FC IDs with loop-attached devices (FL ports) need to remain connected to the same port in which they were configured.

Enabling the Persistent FC ID Feature

To enable the persistent FC ID feature using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx*** and then select **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to enable the Persistent FC ID feature for.

You see the running configuration of the domain in the Information pane.

Step 2 Select the **Persistent Setup** tab and check the **enable** check box for each switch in the fabric that will have persistent FC ID enabled.

Step 3 Click **Apply Changes** to save these changes or click **Undo Changes** to drop any unsaved changes.

About Persistent FC ID Configuration

When the persistent FC ID feature is enabled, you can enter the persistent FC ID submode and add static or dynamic entries in the FC ID database. By default, all added entries are static. Persistent FC IDs are configured on a per-VSAN basis. Follow these requirements to manually configure a persistent FC ID:

- Ensure that the persistent FC ID feature is enabled in the required VSAN.
- Ensure that the required VSAN is an active VSAN—persistent FC IDs can only be configured on active VSANs.
- Verify that the domain part of the FC ID is the same as the runtime domain ID in the required VSAN. If the software detects a domain mismatch, the command is rejected.
- Verify that the port field of the FC ID is 0 (zero) when configuring an area.

Note FICON uses a different scheme for allocating FC IDs based in the front panel port number. This scheme takes precedence over FC ID persistence in FICON VSANs.

Configuring Persistent FC IDs

To configure persistent FC IDs using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx*** and then select **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to configure the Persistent FC ID list for.

You see the running configuration of the domain in the Information pane.

Step 2 Select the **Persistent FcIds** tab and click **Create Row**.

You see the Create Persistent FC IDs dialog box shown in Figure 22-10.

Figure 22-10 Create Persistent FC IDs Dialog Box

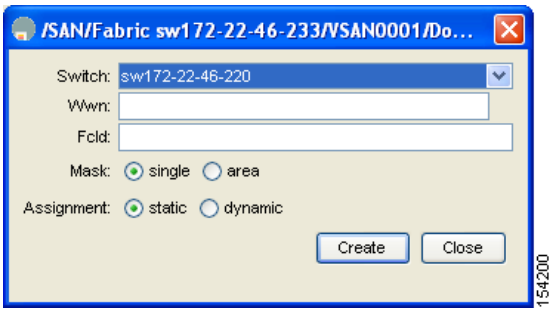

Step 3 Select the switch, WWN, and FC ID that you want to make persistent.

Step 4 Set the Mask radio button to **single** or **area**.

Step 5 Set the Assignment radio button to **static** or **dynamic**.

Step 6 Click **Apply Changes** to save these changes or click **Undo Changes** to drop any unsaved changes.

About Unique Area FC IDs for HBAs

Note Only read this section if the HBA port and the storage port are connected to the same switch.

Some HBA ports require a different area ID than storage ports when they are both connected to the same switch. For example, if the storage port FC ID is 0x6f7704, the area for this port is 77. In this case, the HBA port's area can be anything other than 77. The HBA port's FC ID must be manually configured to be different from the storage port's FC ID.

Switches in the Cisco MDS 9000 Family facilitate this requirement with the FC ID persistence feature. You can use this feature to preassign an FC ID with a different area to either the storage port or the HBA port. The procedure in this example uses a switch domain of 111(6f hex). The HBA port connects to interface fc1/9 and the storage port connects to interface fc 1/10 in the same switch.

Configuring Unique Area FC IDs for an HBA

To configure a different area ID for the HBA port, follow these steps:

Step 1 Expand **End Device** in the Physical Attributes pane and select the **FLOGI** tab in the Information pane to obtain the port WWN (Port Name field) of the HBA. (See Figure 22-11.)

Figure 22-11 FLOGI Database Information in Fabric Manager

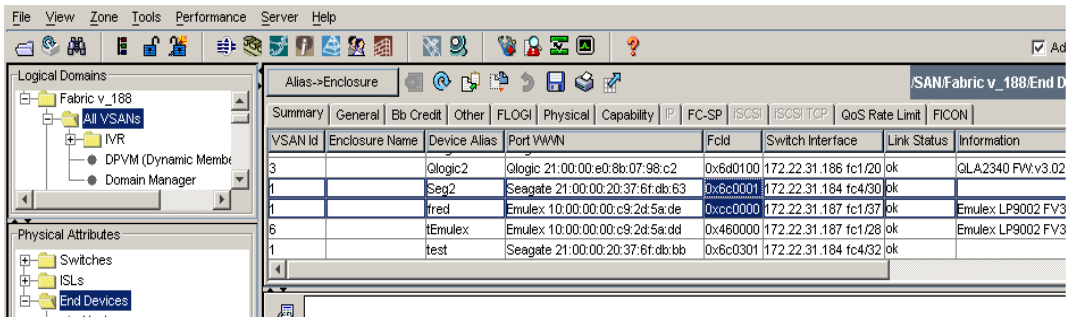

Note Both FC IDs in this setup have the same area 00 assignment.

Step 2 Expand **Switches > Interfaces** and then select **FC Physical** from the Physical Attributes pane.

Step 3 Set the Status Admin drop-down menu to **down** for the interface that the HBA is connected to.

This shuts down the HBA interface in the MDS switch.

Step 4 Expand **Fabricxx > VSANxx** and then select **Domain Manager.**

Step 5 Click the **Persistent Setup** tab in the Information pane to verify that the FC ID feature is enabled. See Figure 22-12.

Figure 22-12 Persistent FC ID Information in Fabric Manager

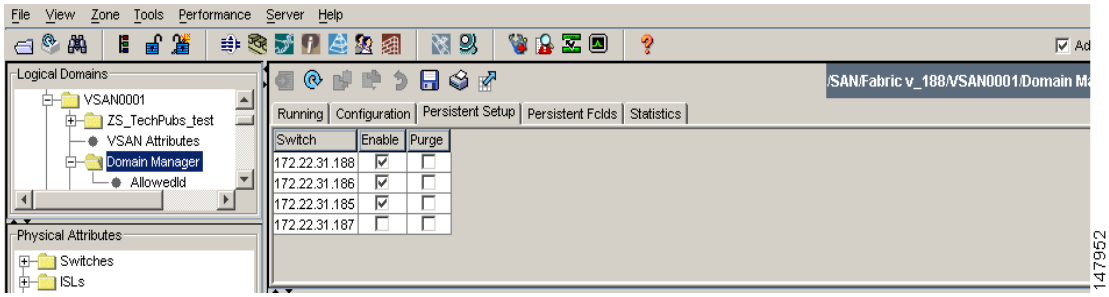

If this feature is disabled, continue with this procedure to enable persistent FC ID.

If this feature is already enabled, skip to Step 7.

Step 6 Check the **Enable** check box to enable the persistent FC ID feature in the Cisco MDS switch. (See Figure 22-13.)

Step 7 Select the **Persistent FcIds** tab and assign a new FC ID with a different area allocation in the FcId field. In this example, we replace *00* with *ee*. (See Figure 22-13.)

Figure 22-13 Setting the FC ID in Fabric Manager

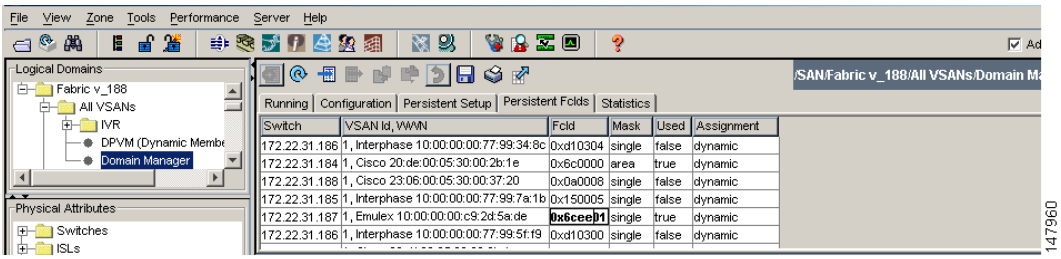

Step 8 Click **Apply Changes** to save this new FC ID.

Step 9 Compare the FC ID values to verify the FC ID of the HBA .

Note Both FC IDs now have different area assignments.

Step 10 Expand **Switches > Interfaces** and then select **FC Physical** from the Physical Attributes pane. Set the Status Admin drop-down menu to **up** for the interface that the HBA is connected to.

This enables the HBA interface in the MDS switch.

About Persistent FC ID Selective Purging

Persistent FC IDs can be purged selectively. Static entries and FC IDs currently in use cannot be deleted. Table 22-1 identifies the FC ID entries that are deleted or retained when persistent FC IDs are purged.

Table 22-1 Purged FC IDs

Persistent FC ID state	Persistent Usage State	Action
Static	In use	Not deleted
Static	Not in use	Not deleted
Dynamic	In use	Not deleted
Dynamic	Not in use	Deleted

Purging Persistent FC IDs

To purge persistent FC IDs using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > All VSANs > Domain Manager** in the Logical Domains pane for the fabric that you want to purge the Persistent FC IDs for. You see the running configuration of the domain in the Information pane.

Step 2 Click the **Persistent Setup** tab.

You see the persistent FC ID setup in the Information pane shown in Figure 22-14.

Figure 22-14 Persistent FC ID Information in Fabric Manager

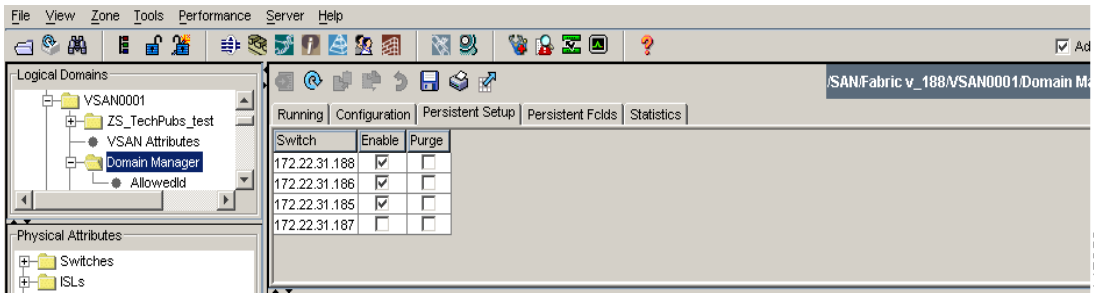

Step 3 Check the **Purge** check box for the switch that you want to purge persistent FC IDs on (see Figure 22-14).

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to drop any unsaved changes.

Displaying fcdomain Statistics

Fabric Manager collects statistics for fcdomain and displays them in the Information pane.

To display fcdomain statistics using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > All VSANs** and then select **Domain Manager** in the Logical Domains pane for the fabric that you want to display statistics for.

You see the running configuration of the domain in the Information pane.

Step 2 Click the **Statistics** tab. You see the FC ID statistics in the Information pane.

Default Settings

Table 22-2 lists the default settings for all fcdomain parameters.

Table 22-2 Default fcdomain Parameters

Parameters	Default
fcdomain feature	Enabled.
Configured domain ID	0 (zero).
Configured domain	Preferred.
autoreconfigure option	Disabled.
contiguous-allocation option	Disabled.
Priority	128.
Allowed list	1 to 239.
Fabric name	20:01:00:05:30:00:28:df.

Table 22-2 Default fcdomain Parameters (continued)

Parameters	Default
rcf-reject	Disabled.
Persistent FC ID	Enabled (as of Release 2.0(1b) this is only configurable on a per-VSAN basis).

Cisco Systems

Send documentation comments to mdsfeedback-doc@cisco.com

PART 4

Fabric Configuration

CHAPTER 23

Configuring and Managing VSANs

You can achieve higher security and greater stability in Fibre Channel fabrics by using virtual SANs (VSANs). VSANs provide isolation among devices that are physically connected to the same fabric. With VSANs you can create multiple logical SANs over a common physical infrastructure. Each VSAN can contain up to 239 switches and has an independent address space that allows identical Fibre Channel IDs (FC IDs) to be used simultaneously in different VSANs. This chapter includes the following sections:

About VSANs

A VSAN is a virtual storage area network (SAN). A SAN is a dedicated network that interconnects hosts and storage devices primarily to exchange SCSI traffic. In SANs you use the physical links to make these interconnections. A set of protocols run over the SAN to handle routing, naming, and zoning. You can design multiple SANs with different topologies.

This section describes VSANs and includes the following topics:

VSAN Topologies

With the introduction of VSANs, the network administrator can build a single topology containing switches, links, and one or more VSANs. Each VSAN in this topology has the same behavior and property of a SAN. A VSAN has the following additional features:

- Multiple VSANs can share the same physical topology.
- The same Fibre Channel IDs (FC IDs) can be assigned to a host in another VSAN, thus increasing VSAN scalability.
- Every instance of a VSAN runs all required protocols such as FSPF, domain manager, and zoning.

- Fabric-related configurations in one VSAN do not affect the associated traffic in another VSAN.
- Events causing traffic disruptions in one VSAN are contained within that VSAN and are not propagated to other VSANs.

As displayed in both Figure 23-1 and Figure 23-2, the switch icons indicate that these features apply to any switch in the Cisco MDS 9000 Family.

Figure 23-1 shows a fabric with three switches, one on each floor. The geographic location of the switches and the attached devices is independent of their segmentation into logical VSANs. Between VSANs no communication is possible. Within each VSAN, all members can talk to one another.

Figure 23-1 Logical VSAN Segmentation

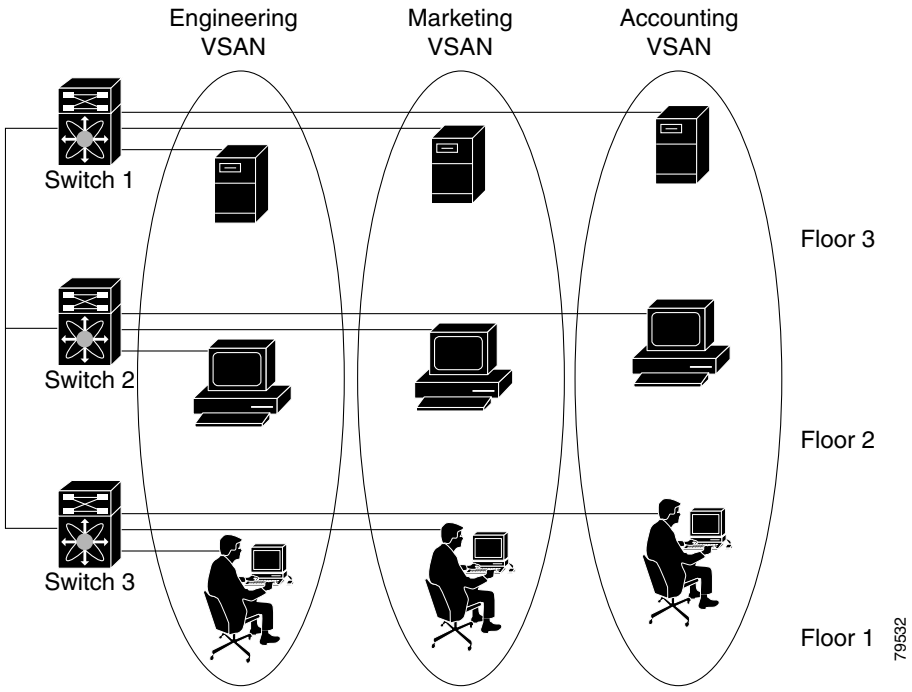

Figure 23-2 shows a physical Fibre Channel switching infrastructure with two defined VSANs: VSAN 2 (dashed) and VSAN 7 (solid). VSAN 2 includes hosts H1 and H2, application servers AS2 and AS3, and storage arrays SA1 and SA4. VSAN 7 connects H3, AS1, SA2, and SA3.

Figure 23-2 Example of two VSANs

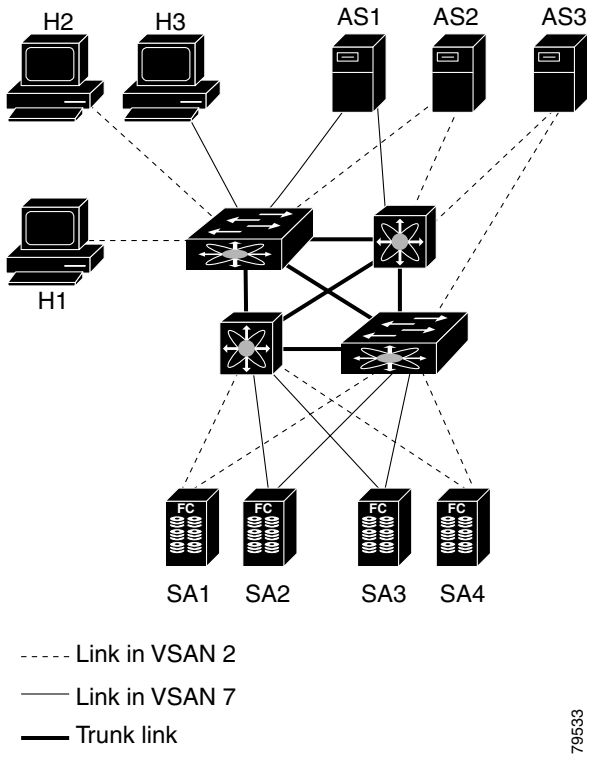

The four switches in this network are interconnected by trunk links that carry both VSAN 2 and VSAN 7 traffic. Thus the inter-switch topology of both VSAN 2 and VSAN 7 are identical. This is not a requirement and a network administrator can enable certain VSANs on certain links to create different VSAN topologies.

Without VSANs, a network administrator would need separate switches and links for separate SANs. By enabling VSANs, the same switches and links may be shared by multiple VSANs. VSANs allow SANs to be built on port granularity instead of switch granularity. Figure 23-2 illustrates that a VSAN is a group of hosts or storage devices that communicate with each other using a virtual topology defined on the physical SAN.

The criteria for creating such groups differ based on the VSAN topology:

- VSANs can separate traffic based on the following requirements:
  - Different customers in storage provider data centers
  - Production or test in an enterprise network
  - Low and high security requirements
  - Backup traffic on separate VSANs
  - Replicating data from user traffic
- VSANs can meet the needs of a particular department or application.

VSAN Advantages

VSANs offer the following advantages:

- Traffic isolation—Traffic is contained within VSAN boundaries and devices reside only in one VSAN ensuring absolute separation between user groups, if desired.
- Scalability—VSANs are overlaid on top of a single physical fabric. The ability to create several logical VSAN layers increases the scalability of the SAN.
- Per VSAN fabric services—Replication of fabric services on a per VSAN basis provides increased scalability and availability.
- Redundancy—Several VSANs created on the same physical SAN ensure redundancy. If one VSAN fails, redundant protection (to another VSAN in the same physical SAN) is configured using a backup path between the host and the device.
- Ease of configuration—Users can be added, moved, or changed between VSANs without changing the physical structure of a SAN. Moving a device from one VSAN to another only requires configuration at the port level, not at a physical level.

Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

VSANs Versus Zones

You can define multiple zones in a VSAN. Because two VSANs are equivalent to two unconnected SANs, zone A on VSAN 1 is different and separate from zone A in VSAN 2. Table 23-1 lists the differences between VSANs and zones.

Table 23-1 VSAN and Zone Comparison

VSAN Characteristic	Zone Characteristic
VSANs equal SANs with routing, naming, and zoning protocols.	Routing, naming, and zoning protocols are not available on a per-zone basis.
—	Zones are always contained within a VSAN. Zones never span two VSANs.
VSANs limit unicast, multicast, and broadcast traffic.	Zones limit unicast traffic.
Membership is typically defined using the VSAN ID to Fx ports.	Membership is typically defined by the pWWN.
An HBA or a storage device can belong only to a single VSAN—the VSAN associated with the Fx port.	An HBA or storage device can belong to multiple zones.
VSANs enforce membership at each E port, source port, and destination port.	Zones enforce membership only at the source and destination ports.
VSANs are defined for larger environments (storage service providers).	Zones are defined for a set of initiators and targets not visible outside the zone.
VSANs encompass the entire fabric.	Zones are configured at the fabric edge.

Figure 23-3 shows the possible relationships between VSANs and zones. In VSAN 2, three zones are defined: zone A, zone B, and zone C. Zone C overlaps both zone A and zone B as permitted by Fibre Channel standards. In VSAN 7, two zones are defined: zone A and zone D. No zone crosses the VSAN boundary—they are completely contained within the VSAN. Zone A defined in VSAN 2 is different and separate from zone A defined in VSAN 7.

Figure 23-3 VSANS with Zoning

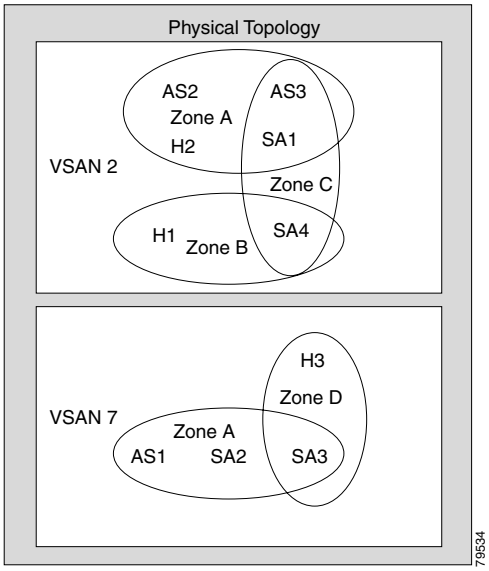

VSAN Configuration

VSANs have the following attributes:

- VSAN ID—The VSAN ID identifies the VSAN as the default VSAN (VSAN 1), user-defined VSANs (VSAN 2 to 4093), and the isolated VSAN (VSAN 4094).
- State—The administrative state of a VSAN can be configured to an active (default) or suspended state. Once VSANs are created, they may exist in various conditions or states.
  - The active state of a VSAN indicates that the VSAN is configured and enabled. By enabling a VSAN, you activate the services for that VSAN.
  - The suspended state of a VSAN indicates that the VSAN is configured but not enabled. If a port is configured in this VSAN, it is disabled. Use this state to deactivate a VSAN without losing the VSAN's configuration. All ports in a suspended VSAN are disabled. By suspending a VSAN, you can preconfigure all the VSAN parameters for the whole fabric and activate the VSAN immediately.
- VSAN name—This text string identifies the VSAN for management purposes. The name can be from 1 to 32 characters long and it must be unique across all VSANs. By default, the VSAN name is a concatenation of VSAN and a four-digit string representing the VSAN ID. For example, the default name for VSAN 3 is VSAN0003.

Note A VSAN name must be unique.

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

About VSAN Creation

A VSAN is in the operational state if the VSAN is active and at least one port is up. This state indicates that traffic can pass through this VSAN. This state cannot be configured.

Creating VSANs Statically

To create and configure VSANs using Fabric Manager, follow these steps:

Step 1 Click the **Create VSAN** icon (see Figure 23-4).

Figure 23-4 Create VSAN Icon

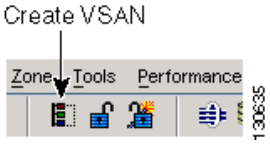

You see the Create VSAN dialog box in Figure 23-5.

Figure 23-5 Create VSAN Dialog Box

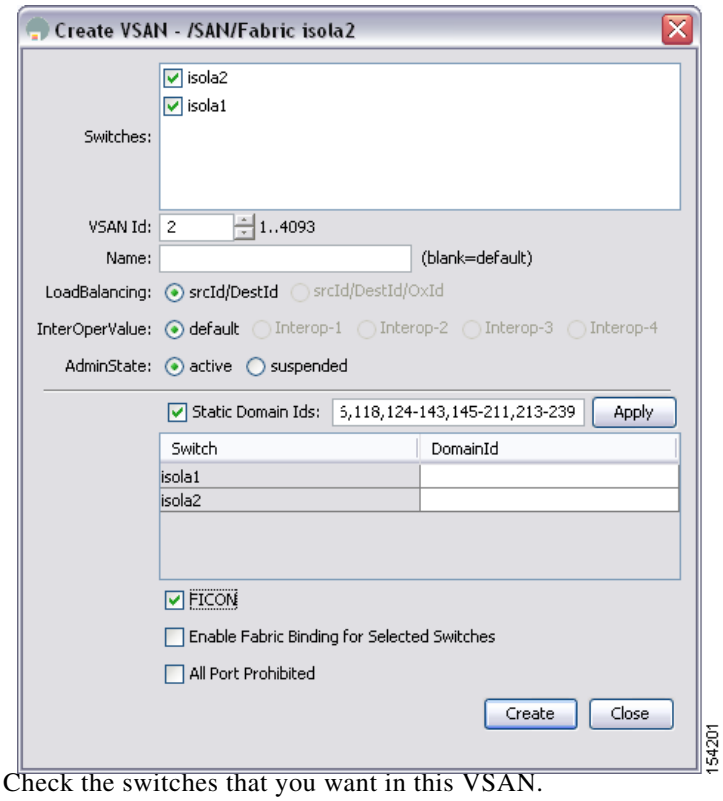

Step 2 Check the switches that you want in this VSAN.

Step 3 Fill in the VSAN Name and VSAN ID fields.

Step 4 Set the LoadBalancing value and the InterOperValue.

Step 5 Set the Admin State to **active** or **suspended**.

Step 6 Check the **Static Domain Ids** check box to assign an unused static domain ID to the VSAN.

Step 7 Optionally, select the **FICON** and **Enable Fabric Binding for Selected Switches** options if you want these features enabled. See Configuring FICON, page 31-1 and Configuring Fabric Binding, page 42-27 for details.

Step 8 Complete the fields in this dialog box and click **Create** to add the VSAN or click **Close**.

About VSAN Membership

Port VSAN membership on the switch is assigned on a port-by-port basis. By default each port belongs to the default VSAN. You can assign VSAN membership to ports using one of two methods:

- Statically—by assigning VSANs to ports.

  For information about changing VSAN membership, see the "Creating VSANs Statically" section on page 23-6.

- Dynamically—by assigning VSANs based on the device WWN. This method is referred to as dynamic port VSAN membership (DPVM).

  See Chapter 24, "Creating Dynamic VSANs."

Trunking ports have an associated list of VSANs that are part of an allowed list (see Chapter 20, "Configuring Trunking").

Assigning Static Port VSAN Membership

To statically assign VSAN membership for an interface using Fabric Manager, follow these steps:

Step 1 Choose **Interfaces > FC Physical** from the Physical Attributes pane. You see the interface configuration in the Information pane.

Step 2 Choose the **General** tab.

You see the Fibre Channel general physical information. Double-click and complete the PortVSAN field.

Step 3 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

About the Default VSAN

The factory settings for switches in the Cisco MDS 9000 Family have only the default VSAN 1 enabled. If you do not need more than one VSAN for a switch, use this default VSAN as the implicit parameter during configuration. If no VSANs are configured, all devices in the fabric are considered part of the default VSAN. By default, all ports are assigned to the default VSAN.

Note VSAN 1 cannot be deleted, but it can be suspended.

Note Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

About the Isolated VSAN

VSAN 4094 is an isolated VSAN. All non-trunking ports are transferred to this VSAN when the VSAN to which they belong is deleted. This avoids an implicit transfer of ports to the default VSAN or to another configured VSAN. All ports in the deleted VSAN are isolated (disabled).

Note When you configure a port in VSAN 4094 or move a port to VSAN 4094, that port is immediately isolated.

Caution Do not use an isolated VSAN to configure ports.

Note Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

Displaying Isolated VSAN Membership

To display interfaces that exist in the isolated VSAN using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*** and then select **All VSANs** in the Logical Domains pane.

You see the VSAN configuration in the Information pane.

Step 2 Click the **Isolated Interfaces** tab.

You see the interfaces that are in the isolated VSAN.

About Static VSAN Deletion

When an active VSAN is deleted, all of its attributes are removed from the running configuration. VSAN-related information is maintained by the system software as follows:

- VSAN attributes and port membership details are maintained by the VSAN manager. This feature is affected when you delete a VSAN from the configuration. When a VSAN is deleted, all the ports in that VSAN are made inactive and the ports are moved to the isolated VSAN. If the same VSAN is recreated, the ports do not automatically get assigned to that VSAN. You must explicitly reconfigure the port VSAN membership (see Figure 23-6).

Figure 23-6 VSAN Port Membership Details

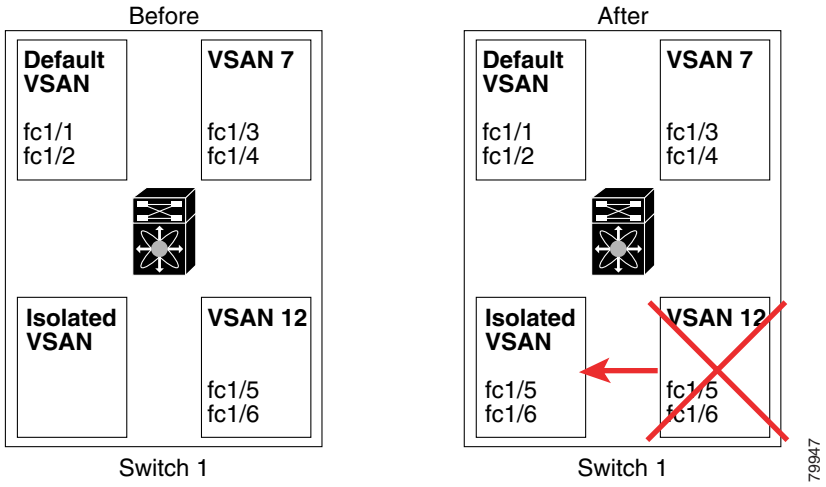

- VSAN-based runtime (name server), zoning, and configuration (static routes) information is removed when the VSAN is deleted.
- Configured VSAN interface information is removed when the VSAN is deleted.

Note The allowed VSAN list is not affected when a VSAN is deleted (see Chapter 20, "Configuring Trunking").

Any commands for a nonconfigured VSAN are rejected. For example, if VSAN 10 is not configured in the system, then a command request to move a port to VSAN 10 is rejected.

Deleting Static VSANs

To delete a VSAN and its attributes using Fabric Manager, follow these steps:

Step 1 Select **All VSANs** from the Logical Domains pane.

The VSANs in the fabric are listed in the Information pane.

Step 2 Right-click the VSAN that you want to delete and select **Delete Row** from the drop-down menu as shown in Figure 23-7.

Figure 23-7 Deleting a VSAN

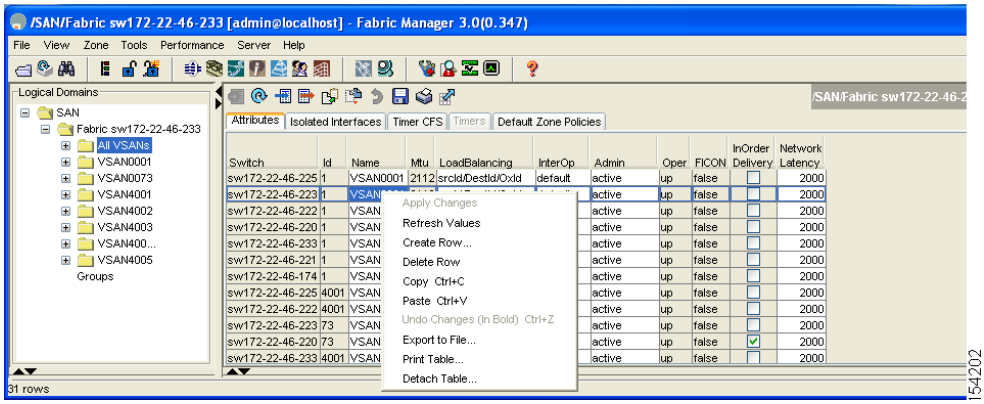

You see a confirmation dialog box.

Step 3 Click **Yes** to confirm the deletion or **No** to close the dialog box without deleting the VSAN.

About Load Balancing

Load balancing attributes indicate the use of the source-destination ID (src-dst-id) or the originator exchange OX ID (src-dst-ox-id, the default) for load balancing path selection.

Configuring Load Balancing

To configure load balancing on an existing VSAN using Fabric Manager, follow these steps:

Step 1 Choose **Fabric*xx* > All VSAN*s*** from the Logical Domains pane.

You see the VSAN configuration in the Information pane shown in Figure 23-8.

Figure 23-8 All VSAN Attributes

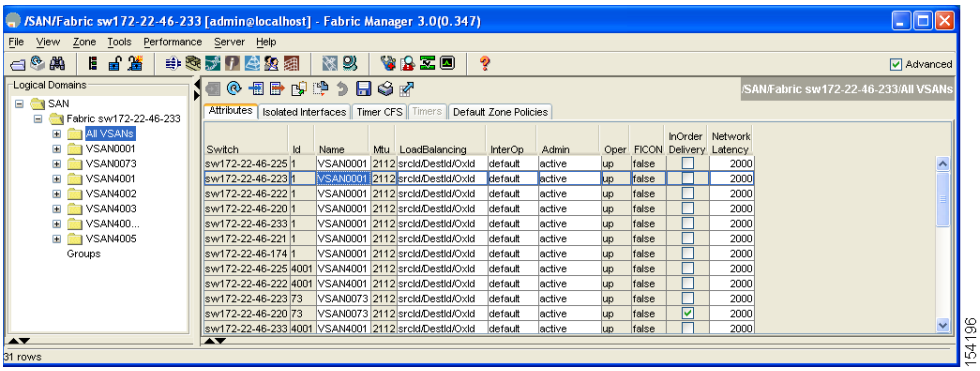

Step 2 Select a VSAN and complete the LoadBalancing field (see Figure 23-8).

Step 3 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

About Interop Mode

Interoperability enables the products of multiple vendors to come into contact with each other. Fibre Channel standards guide vendors towards common external Fibre Channel interfaces. See the "Switch Interoperability" section on page 32-17.

About FICON VSANs

You can enable FICON in up to eight VSANs. See the "FICON VSAN Prerequisites" section on page 31-6.

Default Settings

Table 23-2 lists the default settings for all configured VSANs.

Table 23-2 Default VSAN Parameters

Parameters	Default
Default VSAN	VSAN 1.
State	Active state.
Name	Concatenation of VSAN and a four-digit string representing the VSAN ID. For example, VSAN 3 is VSAN0003.
Load-balancing attribute	OX ID (src-dst-ox-id).

CHAPTER 24

Creating Dynamic VSANs

Port VSAN membership on the switch is assigned on a port-by-port basis. By default each port belongs to the default VSAN.

You can dynamically assign VSAN membership to ports by assigning VSANs based on the device WWN. This method is referred to as Dynamic Port VSAN Membership (DPVM). DPVM offers flexibility and eliminates the need to reconfigure the port VSAN membership to maintain fabric topology when a host or storage device connection is moved between two Cisco MDS switches or two ports within a switch. It retains the configured VSAN regardless of where a device is connected or moved. To assign VSANs statically, see Chapter 23, “Configuring and Managing VSANs.”.

This chapter includes the following sections:

DPVM

DPVM configurations are based on port world wide name (pWWN) and node world wide name (nWWN) assignments. A DPVM database contains mapping information for each device pWWN/nWWN assignment and the corresponding VSAN. The Cisco SAN-OS software checks the database during a device FLOGI and obtains the required VSAN details.

The pWWN identifies the host or device and the nWWN identifies a node consisting of multiple devices. You can assign any one of these identifiers or any combination of these identifiers to configure DPVM mapping. If you assign a combination, then preference is given to the pWWN.

DPVM uses the Cisco Fabric Services (CFS) infrastructure to allow efficient database management and distribution. DPVM uses the application driven, coordinated distribution mode and the fabric-wide distribution scope (see Chapter 12, “Using the CFS Infrastructure”).

Note DPVM does not cause any changes to device addressing. DPVM only pertains to the VSAN membership of the device, ensuring that the host gets same VSAN membership on any port on the switch. For example, if a port on the switch has a hardware failure, you can move the host connection to another port on the switch and not need to update the VSAN membership manually.

Note DPVM is not supported on FL ports. DPVM is supported only on F ports.

This section describes DPVM and includes the following topics:

About DPVM Configuration

To use DPVM as designed, be sure to verify the following requirements:

- The interface through which the dynamic device connects to the Cisco MDS 9000 Family switch must be configured as an F port.
- The static port VSAN of the F port should be valid (not isolated, not suspended and in existence).
- The dynamic VSAN configured for the device in the DPVM database should be valid (not isolated, not suspended and in existence).

Note DPVM overrides any existing static port VSAN membership configuration. If the VSAN corresponding to the dynamic port is deleted or suspended, the port is shut down.

Configuring DPVM with the DPVM Wizard

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

To use the DPVM Setup Wizard in Fabric Manager to set up dynamic port VSAN membership, follow these steps:

Step 1 Click the **DPVM Setup Wizard** icon in the Fabric Manager toolbar. (See Figure 24-1.)

Figure 24-1 DPVM Wizard Icon

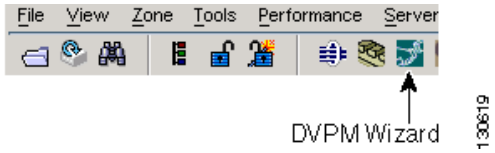

You see the Select Master Switch page.

Step 2 Click the switch you want to be the master switch. This switch controls the distribution of the DPVM database to other switches in the fabric.

Step 3 Click **Next**.

You see the AutoLearn Current End Devices page.

Step 4 Optionally, click the **Create Configuration From Currently Logged In End Devices** check box if you want to turn on autolearning.

Step 5 Click **Next**.

You see the Edit and Activate Configuration page.

Step 6 Verify the current or autolearned configuration. Optionally, click **Insert** to add more entries into the DPVM config database.

Step 7 Click **Finish** to update the DPVM config database, distribute the changes using CFS, and activate the database, or click **Cancel** to exit the DPVM Setup Wizard without saving changes.

About DPVM Databases

The DPVM database consists of a series of device mapping entries. Each entry consists of a device pWWN/nWWN assignment along with the dynamic VSAN to be assigned. You can configure a maximum of 16,000 DPVM entries in the DPVM database. This database is global to the whole switch (and fabric) and is not maintained for each VSAN.

DPVM uses three databases to accept and implement configurations.

- Configuration (config) database—All configuration changes are stored in the configuration database when distribution is disabled.
- Active database—The database currently enforced by the fabric.
- Pending database—All configuration changes are stored in the pending database when distribution is enabled (see the "About DPVM Database Distribution" section on page 24-11).

Changes to the DVPM config database are not reflected in the DVPM active database until you activate the DVPM config database. Changes to the pending database are not reflected in the config/active database until you commit the pending database. This database structure allows you to create multiple entries, review changes, and let the config and pending databases take effect.

Figure 24-2 shows an example of the DPVM databases in the Information pane in Fabric Manager.

Figure 24-2 DPVM Configuration in Fabric Manager

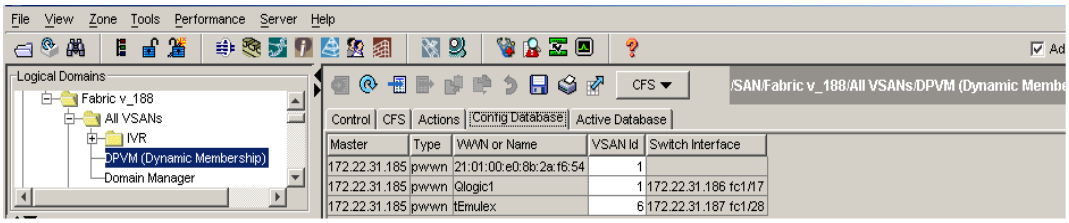

Configuring Config and Pending Databases

To create and populate the config and pending databases using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*> All VSANs** and then select **DPVM** in the Logical Attributes pane.

You see the DPVM configuration in the Information pane.

Step 2 Click the **CFS** tab and select a master switch by checking a check box in the Master column (see Figure 24-3).

Figure 24-3 CFS Tab with Master Switch Checked

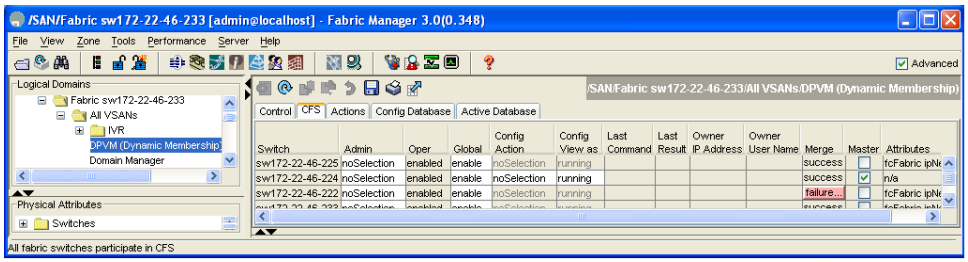

Step 3 Click the **Config Database** tab and then click the **Create Row** to insert a new entry.

You see the Create Config Database dialog box shown in Figure 24-4.

Figure 24-4 Create Config Database

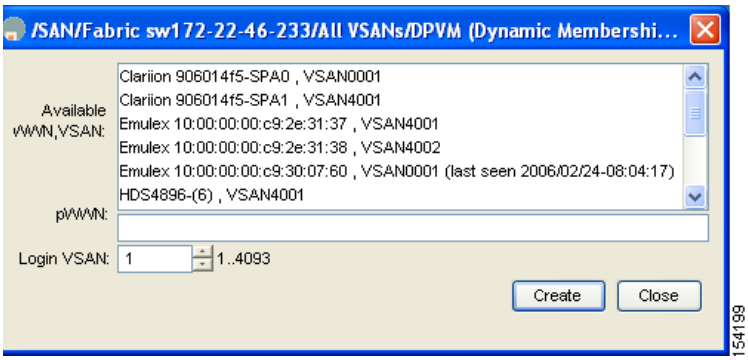

Step 4 Choose an available WWN and VSAN combination or fill in the pWWN and Login VSAN fields.

Step 5 Click **Create** to save these changes in the config or pending database or click **Close** to discard any unsaved changes.

Step 6 Click the **CFS** tab and select the Config Action drop-down menu for the master database.

You see the options shown in Figure 24-5.

Figure 24-5 Config Action Drop-down Menu

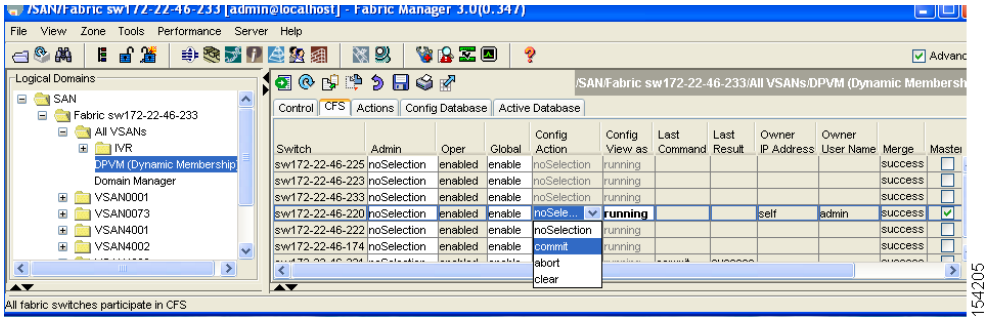

Step 7 Select **commit** from the drop-down menu to distribute these changes or **abort** to discard the changes (see Figure 24-5).

Activating Config Databases

When you explicitly activate the DVPM config database, the DVPM config database becomes the DVPM active database. Activation may fail if conflicting entries are found between the DVPM config database and the currently active database. However, you can force activation to override conflicting entries.

To activate the DVPM config database using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*> All VSANs** and then select **DPVM** from the Logical Attributes pane.

You see the DPVM configuration in the Information pane.

Step 2 Click the **Action** tab and set the Action drop-down menu to **activate** or **forceActivate** to activate the DVPM config database (see Figure 24-6).

Figure 24-6 Activate a Configured Database

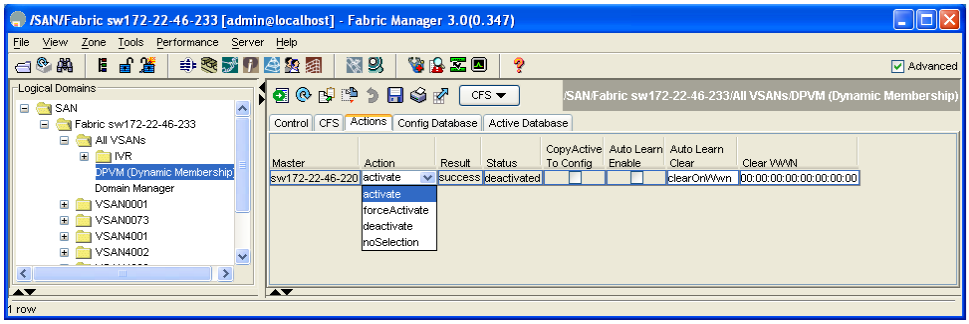

Step 3 Click the **CFS** tab and select the Config Action drop-down menu for the master database.

You see the options shown in Figure 24-7.

Figure 24-7 Config Action Drop-down Menu

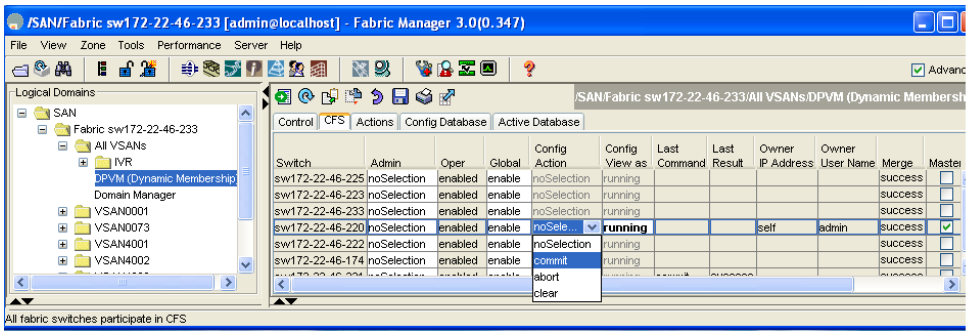

Step 4 Select **commit** from the drop-down menu to distribute these changes or **abort** to discard the changes.

Note To disable DPVM, you must explicitly deactivate the currently active DPVM database.

Viewing the Pending Database

To view the pending database using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*> All VSANs** and then select **DPVM** from the Logical Attributes pane.

You see the DPVM configuration in the Information pane.

Step 2 Click the **CFS** tab (see Figure 24-8) and set the Config View drop-down menu to **pending**.

Figure 24-8 CFS Tab with Master Switch Checked

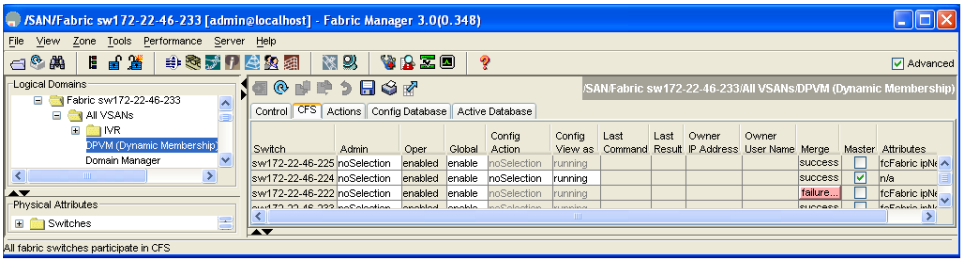

Step 3 Click **Apply Changes**.

Step 4 Click the **Config Database** tab.

You see the pending database entries.

About Autolearned Entries

The DPVM database can be configured to automatically learn (autolearn) about new devices within each VSAN. Autolearn can be enabled or disabled at any time. Learned entries are created by populating device pWWNs and VSANs in the DPVM active database. The DPVM active database should already be available to enable autolearn.

You can delete any learned entry from the DPVM active database when you enable autolearn. These entries only become permanent in the DVPM active database when you disable autolearn.

The following conditions apply to learned entries:

- If a device logs out while autolearn is enabled, that entry is automatically deleted from the DPVM active database.
- If the same device logs multiple times into the switch through different ports, then the VSAN corresponding to last login is remembered.
- Learned entries do not override previously configured and activated entries.
- Learning is a two-part process—enabling autolearning followed by disabling autolearning.When the autolearn option is enabled, the following applies:
  - Learning currently logged-in devices—occurs from the time learning is enabled.
  - Learning new device logins— occurs as and when new devices log in to the switch.

Enabling Autolearning

To enable autolearning using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*> All VSANs** and then select DPVM from the Logical Attributes pane.

You see the DPVM configuration in the Information pane.

Step 2 Click the **Actions** tab and check the **Auto Learn Enable** check box to enable autolearning (see Figure 24-9).

Figure 24-9 DPVM Actions Tab

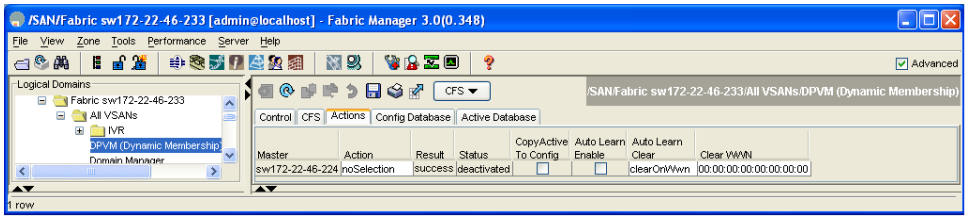

Step 3 Click the **CFS** tab and select **commit** to distribute these changes or **abort** to discard the changes.

Clearing Learned Entries

You can clear DPVM entries from the DVPM active database (if autolearn is still enabled) using one of two methods

To clear a single autolearn entry using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*> All VSANs** and then select **DPVM** from the Logical Attributes pane.

You see the DPVM configuration in the Information pane.

Step 2 Click the **Actions** tab and select **clearOnWWN** from the Auto Learn Clear drop-down men.

Step 3 Check the **clear WWN** check box next to the WWN of the autolearned entry that you want to clear.

Step 4 Click **CFS** and select **commit** to distribute these changes or **abort** to discard the changes.

To clear all autolearn entries using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*> All VSANs** and then select **DPVM** from the Logical Attributes pane.

You see the DPVM configuration in the Information pane.

Step 2 Click the **Actions** tab.

You see the DPVM Actions menu shown in Figure 24-10.

Figure 24-10 DPVM Actions Tab

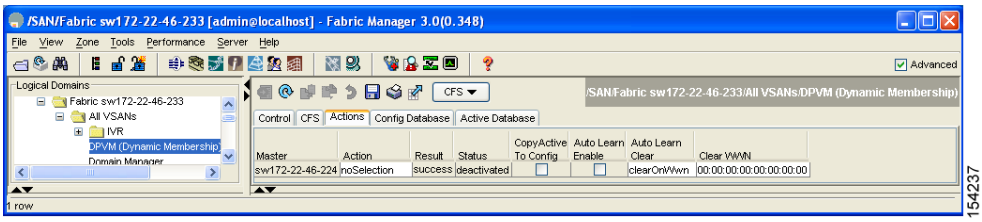

Step 3 Select **clear** from the Auto Learn Clear drop-down menu.

Step 4 Click the **CFS** tab and select **commit** to distribute these changes or **abort** to discard the changes.

Note These two procedures do not start a session and can only be issued in the local switch.

DPVM Database Distribution

If the DPVM database is available on all switches in the fabric, devices can be moved anywhere and offer the greatest flexibility. To enable database distribution to the neighboring switches, the database should be consistently administered and distributed across all switches in the fabric. The Cisco SAN-OS software uses the Cisco Fabric Services (CFS) infrastructure to achieve this requirement (see Chapter 12, "Using the CFS Infrastructure").

This section describes how to distribute the DPVM database and includes the following topics:

About DPVM Database Distribution

Using the CFS infrastructure, each DPVM server learns the DPVM database from each of its neighboring switches during the ISL bring-up process. If you change the database locally, the DPVM server notifies its neighboring switches, and that database is updated by all switches in the fabric.

If fabric distribution is enabled, all changes to the configuration database are stored in the pending database. These changes include the following tasks:

- Adding, deleting, or modifying database entries.
- Activating, deactivating, or deleting the configuration database.
- Enabling or disabling autolearning.

These changes are distributed to all switches in a fabric when you commit the changes. You can also discard (abort) the changes at this point.

Tip See the "Viewing the Pending Database" section on page 24-8 to view the contents of the pending database.

Disabling DPVM Database Distribution

To disable DPVM database distribution to the neighboring switches using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*> All VSANs** and then select **DPVM** from the Logical Attributes pane.

You see the DPVM configuration in the Information pane.

Step 2 Click the **CFS** tab and select **disable** from the Admin drop-down menu.

Step 3 Click **Apply Changes** to save this change or click **Undo Changes** to discard the change.

About Locking the Fabric

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

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database. Modifications from this point on are made to the pending database. The pending database remains in effect until you commit the modifications to the pending database or discard (abort) the changes to the pending database.

Locking the Fabric

To lock the fabric and apply changes to the pending database using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*> All VSANs** and then select **DPVM** from the Logical Attributes pane.

You see the DPVM configuration in the Information pane.

Step 2 Click the **Config Database** tab and **Create Row**.

You see the DPVM Config Database Create dialog box shown in Figure 24-11.

Figure 24-11 DPVM Create Config Database

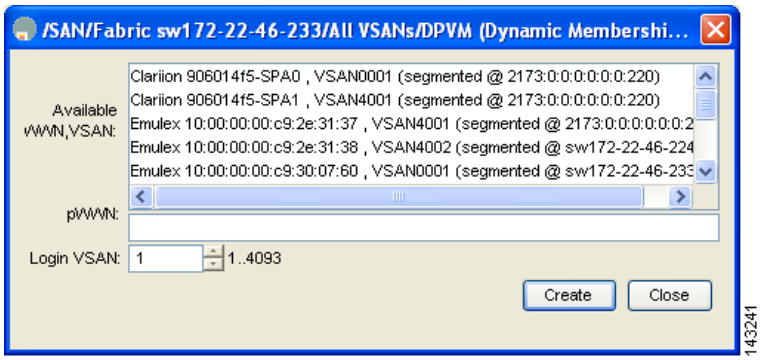

Step 3 Choose an available pWWN and login VSAN (see Figure 24-11).

Step 4 Click **Create** to save this change to the pending database or click **Close** to discard any unsaved change.

Committing Changes

If you commit the changes made to the configuration, the configuration in the pending database are distributed to other switches. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

To commit the pending database using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*> All VSANs** and then select **DPVM** from the Logical Attributes pane.

You see the DPVM configuration in the Information pane.

Step 2 Click the **CFS** tab and select **commit from** the Config Action drop-down menu.

Step 3 Click **Apply Changes** to save this change or click **Undo Changes** to discard the change.

Discarding Changes

If you discard (abort) the changes made to the pending database, the configurations remain unaffected and the lock is released.

To discard the pending database using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*> All VSANs** and then select **DPVM** from the Logical Attributes pane.

You see the DPVM configuration in the Information pane.

Step 2 Click the **CFS** tab and select **abort** from the Config Action drop-down menu.

Step 3 Click **Apply Changes** to save this change or click **Undo Changes** to discard the change.

Clearing a Locked Session

If you have performed a DPVM task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.

Tip The pending database is only available in the volatile directory and are subject to being discarded if the switch is restarted.

To use administrative privileges and release a locked DPVM session using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*> All VSANs** and then select **DPVM** from the Logical Attributes pane.

You see the DPVM configuration in the Information pane.

Step 2 Click the **CFS** tab and select **clear** from the Config Action drop-down menu.

Step 3 Click **Apply Changes** to save this change or click **Undo Changes** to discard the change.

Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the DVPM active database. See the "CFS Merge Support" section on page 12-9 for detailed concepts.

When merging the database between two fabric, follow these guidelines:

- Verify that the activation status and the auto-learn status is the same is both fabrics.
- Verify that the combined number of device entries in each database does not exceed 16K.

Caution If you do not follow these two conditions, the merge will fail. The next distribution will forcefully synchronize the databases and the activation states in the fabric.

This section describes how to merge DPVM databases and includes the following topics:

About Copying DPVM Databases

The following circumstances may require the DVPM active database to be copied to the DVPM config database:

- If the learned entries are only added to the DVPM active database.
- If the DVPM config database or entries in the DVPM config database are accidently deleted.

Note If you copy the DPVM database and fabric distribution is enabled, you must commit the changes.

Copying DPVM Databases

To copy the currently active database to the DVPM config database using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*> All VSANs** and then select **DPVM** in the Logical Attributes pane.

You see the DPVM configuration in the Information pane.

Step 2 Click the **Actions** tab and check the **CopyActive to Config** check box.

Step 3 Click the **CFS** tab and select **commit** from the Config Action drop-down menu.

Comparing Database Differences

To compare the currently active database to the DVPM config database using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*> All VSANs** and then select **DPVM** from the Logical Attributes pane.

You see the DPVM configuration in the Information pane.

Step 2 Click the **Active Database** tab.

You see the DPVM active database in the Information pane.

Step 3 Select **Config** from the Compare With drop-down menu.

You see the comparison dialog box.

Step 4 Select **Close** to close the comparison dialog box.

Default Settings

Table 24-1 lists the default settings for DPVM parameters.

Table 24-1 Default DPVM Parameters

Parameters	Default
DPVM	Disabled.
DPVM distribution	Enabled.
Autolearning	Disabled.

CHAPTER 25

Configuring Inter-VSAN Routing

This chapter explains the inter-VSAN routing (IVR) feature and provides details on sharing resources across VSANs using IVR management interfaces provided in the switch.

This chapter includes the following sections:

Inter-VSAN Routing

Virtual SANs (VSANs) improve storage area network (SAN) scalability, availability, and security by allowing multiple Fibre Channel SANs to share a common physical infrastructure of switches and ISLs. These benefits are derived from the separation of Fibre Channel services in each VSAN and isolation of traffic between VSANs. Data traffic isolation between the VSANs also inherently prevents sharing of resources attached to a VSAN, like robotic tape libraries. Using IVR, resources across VSANs are accessed without compromising other VSAN benefits.

This section includes the following topics:

About IVR

Data traffic is transported between specific initiators and targets on different VSANs without merging VSANs into a single logical fabric. Fibre Channel control traffic does not flow between VSANs, nor can initiators access any resource across VSANs aside from the designated ones. Valuable resources such as tape libraries are easily shared across VSANs without compromise.

IVR is in compliance with Fibre Channel standards and incorporates third-party switches, however, IVR-enabled VSANs may have to be configured in one of the interop modes.

IVR is not limited to VSANs present on a common switch. Routes that traverse one or more VSANs across multiple switches can be established, if necessary, to establish proper interconnections. IVR used in conjunction with FCIP provides more efficient business continuity or disaster recovery solutions (see Figure 25-1).

Figure 25-1 Traffic Continuity Using IVR and FCIP

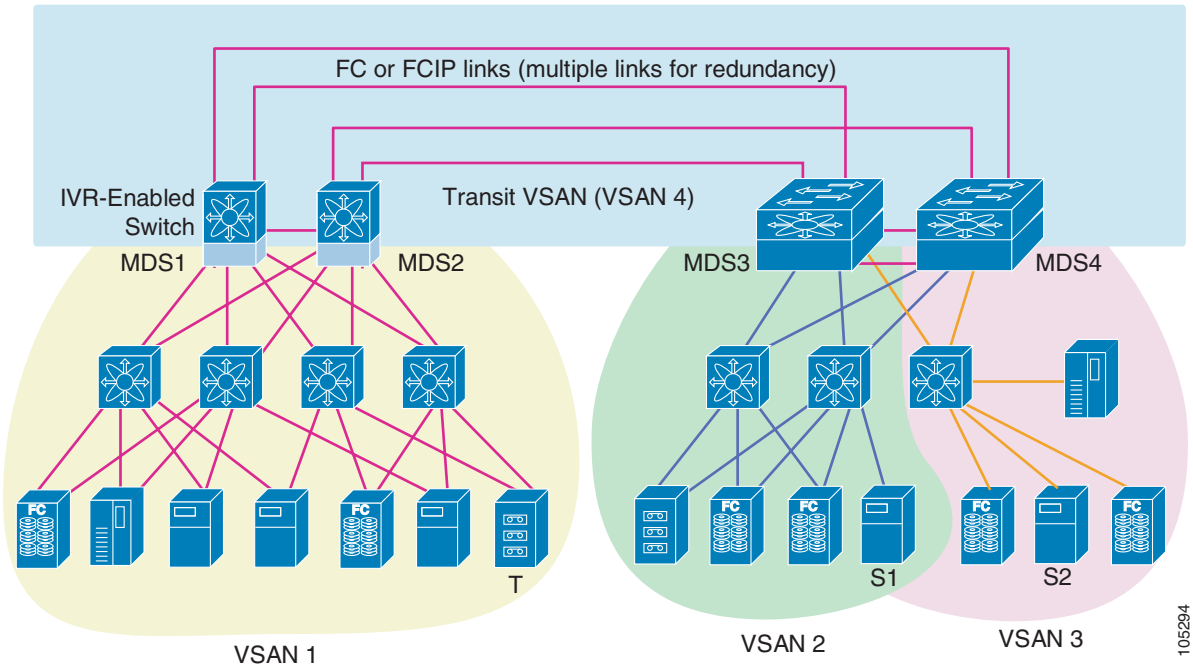

IVR Terminology

The following terms are used in this chapter.

- Native VSAN—The VSAN to which an end device logs on is the native VSAN for that end device.
- Inter-VSAN routing zone (IVR zone)—A set of end devices that are allowed to communicate across VSANs within their interconnected SAN fabric. This definition is based on their port world wide names (pWWNs) and their native VSAN associations. You can configure up to 2,000 IVR zones and10,000 IVR zone members in the fabric from any switch in the Cisco MDS 9000 Family.
- Inter-VSAN routing zone sets (IVR zone sets)—One or more IVR zones make up an IVR zone set. You can configure up to 32 IVR zone sets on any switch in the Cisco MDS 9000 Family. Only one IVR zone set can be active at any time.

- IVR path—An IVR path is a set of switches and Inter-Switch Links (ISLs) through which a frame from one end-device in one VSAN can reach another end-device in some other VSAN. Multiple paths can exist between two such end-devices.
- IVR-enabled switch—A switch in which the IVR feature is enabled.
- Edge VSAN—A VSAN that initiates (source edge-VSAN) or terminates (destination edge-VSAN) an IVR path. Edge VSANs may be adjacent to each other or they may be connected by one or more transit VSANs. In Figure 25-1, VSANs 1, 2, and 3 are edge VSANs.

Note An edge VSAN for one IVR path can be a transit VSAN for another IVR path.

- Transit VSAN—A VSAN that exists along an IVR path from the source edge VSAN of that path to the destination edge VSAN of that path. In Figure 25-1, VSAN 4 is a transit VSAN.

Note When the source and destination edge VSANs are adjacent to each other, then a transit VSAN is not required between them.

- Border switch—An IVR-enabled switch that is a member of two or more VSANs. Border switches in Figure 25-1 span two or more different color-coded VSANs.
- Edge switch—A switch to which a member of an IVR zone has logged in. Edge switches are oblivious to the IVR configurations in the border switches. Edge switches need not be IVR enabled.

Fibre Channel Header Modifications

IVR works by virtualizing the remote end devices in the native VSAN using a virtual domain. When IVR is configured to link end devices in two disparate VSANs, the IVR border switches are responsible for modifying the Fibre Channel headers for all communication between the end devices. The sections of the Fibre Channel frame headers that are modified include:

- VSAN number
- Source FCID
- Destination FCID

When a frame goes from the initiator to the target, the Fibre Channel frame header is modified such that the initiator VSAN number is changed to the target VSAN number. If IVR Network Address Translation (NAT) is enabled, then the source and destination FCIDs are also translated at the edge border switch. If IVR NAT is not enabled, then you must configure unique domain IDs for all switches involved in the IVR path.

IVR NAT

IVR NATallows you to set up IVR in a fabric without needing unique domain IDs on every switch in the IVR path. IVR NAT virtualizes the switches in other VSANs by using local VSAN for the destination IDs in the Fibre Channel headers. In some Extended Link Service message types, the destinations IDs are part of the payload. In these cases, IVR NAT replaces the actual destination ID with the virtualized destination ID. IVR NAT supports destination ID replacement in the Extended Link Service messages described in Table 25-1.

Table 25-1 Extended Link Service Messages Supported by IVR NAT

Extended Link Service Messages	Link Service Command (LS_COMMAND)	Mnemonic
Abort Exchange	0x06 00 00 00	ABTX
Discover Address	0x52 00 00 00	ADISC
Discover Address Accept	0x02 00 00 00	ADISC ACC
Fibre Channel Address Resolution Protocol Reply	0x55 00 00 00	FARP-REPLY
Fibre Channel Address Resolution Protocol Request	0x54 00 00 00	FARP-REQ
Logout	0x05 00 00 00	LOGO
Port Login	0x30 00 00 00	PLOGI
Read Exchange Concise	0x13 00 00 00	REC
Read Exchange Concise Accept	0x02 00 00 00	REC ACC
Read Exchange Status Block	0x08 00 00 00	RES
Read Exchange Status Block Accept	0x02 00 00 00	RES ACC
Read Link Error Status Block	0x0F 00 00 00	RLS
Read Sequence Status Block	0x09 00 00 00	RSS
Reinstate Recovery Qualifier	0x12 00 00 00	RRQ
Request Sequence Initiative	0x0A 00 00 00	RSI
Scan Remote Loop	0x7B 00 00 00	RSL
Third Party Process Logout	0x24 00 00 00	TPRLO
Third Party Process Logout Accept	0x02 00 00 00	TPRLO ACC

If you have a message that is not recognized by IVR NAT and contains the destination ID in the payload, you cannot use IVR with NAT in your topology. You can still use IVR with unique domain IDs.

Note IVR NAT requires Cisco MDS SAN-OS Release 2.1(1a) or later on all switches in the fabric performing IVR. If you have isolated switches with an earlier release that are involved in IVR, you must remove any isolated fabrics from monitoring by Fabric Manager server and then re-open the fabric to use IVR NAT. See the "Removing a Fabric from Monitoring" section on page 3-7.

IVR VSAN Topology

IVR uses a configured IVR VSAN topology to determine how to route traffic between the initiator and the target across the fabric. You can configure this IVR VSAN topology manually on an IVR-enabled switch and distribute the configuration using CFS in Cisco MDS SAN-OS Release 2.0(1b) or later. Alternately, in Cisco MDS SAN-OS Release 2.1(1a) or later, you can configure IVR topology in auto mode. Prior to Cisco MDS SAN-OS Release 2.0(1b), you need to manually copy the IVR VSAN topology to each switch in the fabric.

Auto mode automatically builds the IVR VSAN topology and maintains the topology database when fabric reconfigurations occur. Auto mode distributes the IVR VSAN topology to IVR-enabled switches using CFS.

Using auto mode, you no longer need to manually update the IVR VSAN topology when reconfigurations occur in your fabric. If a manually configured IVR topology database exists, auto mode initially uses that topology information. This reduces disruption in the network by gradually migrating from the user-specified topology database to the automatically learned topology database. User configured topology entries that are not part of the network are aged out in about three minutes. New entries that are not part of the user configured database are added as they are learned from the network.

Note IVR topology in auto mode requires Cisco MDS SAN-OS Release 2.1(1a) or later and enabling CFS for IVR on all switches in the fabric.

Autonomous Fabric ID

The autonomous fabric ID (AFID) distinguishes segmented VSANS (that is, two VSANs that are logically and physically separate but have the same VSAN number). Cisco MDS SAN-OS supports AFIDs from 1 through 64. AFIDs are used in conjunction with auto mode to allow segmented VSANS in the IVR VSAN topology database. You can configure up to 64 AFIDs.

The AFID can be configured individually for each switch and list of VSANs, or the default AFID can be configured for each switch.

Note Two VSANs with the same VSAN number but different AFIDs are counted as two VSANs out of the total 128 VSANs allowed in the fabric.

IVR Interoperability

When using the IVR feature, all border switches in a given fabric must be Cisco MDS switches. However, other switches in the fabric may be non-MDS switches. For example, end devices that are members of the active IVR zone set may be connected to non-MDS switches. Non-MDS switches may also be present in the transit VSAN(s) or in the edge VSANs if one of the **interop** modes is enabled.

See the "Switch Interoperability" section on page 32-17.

About the IVR Zone Wizard

The IVR Zone Wizard simplifies the steps required to configure IVR zones in a fabric. The IVR Zone Wizard checks the following conditions and prompts you for any issues:

- Checks if all switches in the fabric are Cisco MDS SAN-OS Release 2.1(1a) or later and if so, asks if you want to migrate to using IVR NAT with auto-topology.
- Checks if any switches in the fabric are earlier than Cisco MDS SAN-OS Release 2.1(1a) and if so, asks you to upgrade the necessary switches or to disable IVR NAT or auto-topology if they are enabled.

Configuring IVR Using the IVR Zone Wizard

To configure IVR and IVR zones using the IVR Zone Wizard in Fabric Manager, follow these steps:

Step 1 Click the **IVR Zone Wizard** icon in the Zone toolbar. (See Figure 25-2.)

Figure 25-2 IVR Zone Wizard Icon

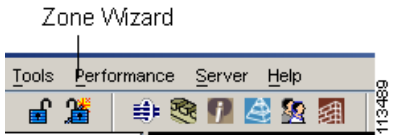

To migrate to IVR NAT mode click **Yes**, otherwise click **No**. You see the IVR Zone Wizard.

Step 2 Select the VSANs that will participate in IVR in the fabric.

Step 3 Select the end devices that you want to communicate over IVR.

Note If you are not using IVR NAT, Fabric Manager may display an error message if all the switches participating in IVR do not have unique domain IDs. You must reconfigure those switches before configuring IVR. Go to Step 5.

Step 4 If you enable IVR NAT, verify switches that Fabric Manager will enable with IVR NAT, CFS for IVR, and IVR topology in auto mode.

Step 5 Enter the VSAN ID of the VSAN you want to use as the transit VSAN between the VSANs selected for the IVR zone. Click **Next**.

Step 6 Optionally, configure a unique AFID for switches in the fabric that have non-unique VSAN IDs in the Select AFID dialog box.

Step 7 If you did not enable IVR NAT, verify the transit VSAN or configure the transit VSAN if Fabric Manager cannot find an appropriate transit VSAN.

Step 8 Set the IVR zone and IVR zone set.

Step 9 Verify all steps that Fabric Manager will take to configure IVR in the fabric.

Step 10 Click **Finish** if you want to enable IVR NAT and IVR topology and to create the associated IVR zones and IVR zone set, or click **Cancel** to exit the IVR Wizard without saving any changes.

Step 11 The Save Configuration dialog is displayed. You can save the configuration of the master switch to be copied to other IVR-enabled switches. Click **Continue Activation** or you may click **Cancel**.

Step 12 Click **Finish**.

Note IVR NAT and auto-topology can be configured independently if you configure these features outside the IVR Zone Wizard. See the "Manual IVR Configuration" section on page 25-7.

Manual IVR Configuration

You can configure IVR using the IVR tables in the Information pane in Fabric Manager. Use these tables only if you are familiar with all IVR concepts. We recommend you configure IVR using the IVR Wizard.

Note Most tabs in the Information pane for features using CFS are dimmed until you click the CFS tab. The CFS tab shows which switches have CFS enabled and shows the master switch for this feature. Once the CFS tab is clicked, the other tabs in the Information pane are activated.

This section describes manually configuring IVR and includes the following topics:

About IVR NAT and Auto Topology

Before configuring an IVR SAN fabric to use IVR NAT and auto-topology, consider the following guidelines:

- Configure IVR only in the relevant switches.
- Enable CFS for IVR on all switches in the fabric.
- Verify that all switches in the fabric are running Cisco MDS SAN-OS Release 2.1(1a) or later.
- Acquire a mandatory Enterprise License Package or SAN-EXTENSION license package if you have Cisco MDS SAN-OS Release2.1(1a) or later and one active IPS card for this feature.

Tip If you change any FSPF link cost, ensure that the FSPF path distance (that is, the sum of the link costs on the path) of any IVR path is less than 30,000.

Note IVR-enabled VSANs can be configured when the interop mode is enabled (any interop mode) or disabled (no interop mode).

Transit VSAN Guidelines

Consider the following guidelines for transit VSANs:

- Besides defining the IVR zone membership, you can choose to specify a set of transit VSANs to provide connectivity between two edge VSANs:
  - If two edge VSANs in an IVR zone overlap, then a transit VSAN is not required (though, not prohibited) to provide connectivity.
  - If two edge VSANs in an IVR zone do not overlap, you may need one or more transit VSANs to provide connectivity. Two edge VSANs in an IVR zone will not overlap if IVR is not enabled on a switch that is a member of both the source and destination edge VSANs.
- Traffic between the edge VSANs only traverses through the shortest IVR path.
- Transit VSAN information is common to all IVR zone sets. Sometimes, a transit VSAN can also act as an edge VSAN in another IVR zone.

Border Switch Guidelines

Before configuring border switches, consider the following guidelines:

- Border switches require Cisco MDS SAN-OS Release 2.1(1a) or later.
- A border switch must be a member of two or more VSANs.
- A border switch that facilities IVR communications must be IVR enabled.
- IVR can (optionally) be enabled on additional border switches to provide redundant paths between active IVR zone members.

The VSAN topology configuration updates automatically when a border switch is added or removed.

Configuring IVR NAT and IVR Auto Topology

To configure IVR in NAT mode and IVR topology in auto mode from Fabric Manager, follow these steps:

Step 1 Expand **All VSANs** and then select **IVR** in the Logical Domains pane.

You see the inter-VSAN routing configuration in the Information pane shown in Figure 25-3.

Figure 25-3 IVR Routing Configuration Control Tab

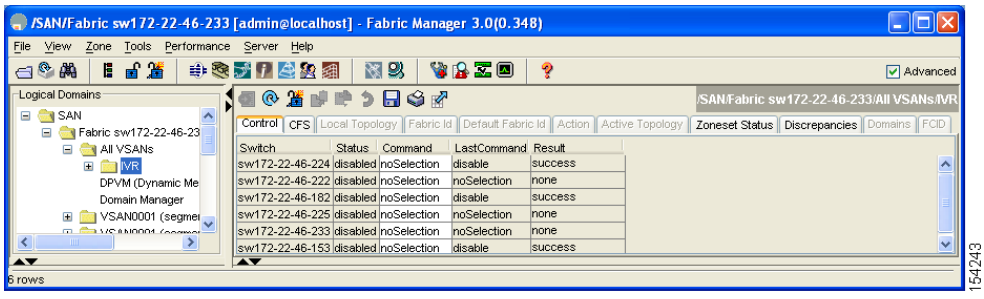

Step 2 Click the **CFS** tab if CFS is enabled for this feature in the fabric.

Step 3 Select **enable** from the Admin column drop-down menu for the primary switch.

Step 4 Click **Apply Changes** from the Information pane to distribute this change to all switches in the fabric, or click **Undo Changes** to cancel any changes you made.

Step 5 Click the **Action** tab.

Step 6 Check the **Enable IVR NAT** check box to enable IVR in NAT mode.

Step 7 Check the **Auto Discover Topology** check box to enable IVR topology in auto mode.

Step 8 Click **Apply Changes** from the Information pane to enable IVR on the switches, or click **Undo Changes** to cancel any changes you made.

Step 9 Click **CFS > Config Changes > Action** and select **commit**.

Step 10 Click **Apply Changes** from the Information pane to distribute IVR on the switches.

About AFIDs

You configure AFIDs individually for VSANs, or you set the default AFIDs for all VSANs on a switch. If you configure an individual AFID for a subset of the VSANs on a switch that has a default AFID, that subset uses the configured AFID while all other VSANs on that switch use the default AFID.

Configuring Default AFIDs

To configure default AFIDs using Fabric Manager, follow these steps:

Step 1 Expand **All VSANs** and then select **IVR** in the Logical Domains pane.

Step 2 You see the IVR configuration in the Information pane shown in Figure 25-4.

Figure 25-4 IVR Routing Configuration Control Tab

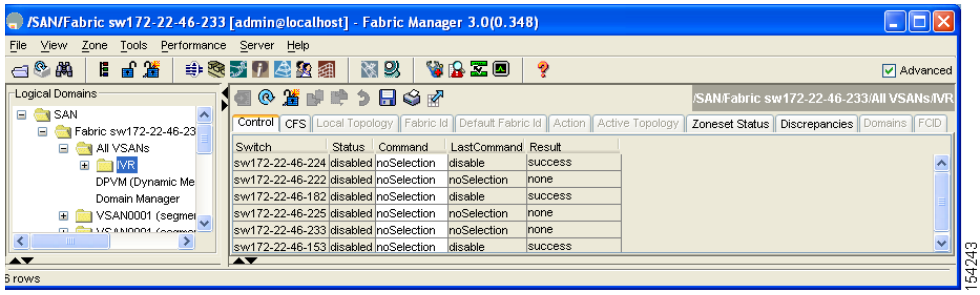

Step 3 Click the **Default Fabric ID** tab to display the existing default AFIDs. (If this tab is greyed out, click the CFS tab first.)

Step 4 Click the **Create Row** icon to create a default AFID.

You see the default AFID dialog box.

Step 5 Check the check boxes next to each switch involved in IVR that you want to use this default AFID.

Step 6 Provide a name for each SwitchWWN and set the default Fabric ID.

Step 7 Click **Create** to create this entry or click **Cancel** to discard all changes.

Step 8 Repeat Step 1 through Step 7 for all default AFIDs that you want to configure in your IVR topology.

Configuring Individual AFIDs

To configure individual AFIDs using Fabric Manager, follow these steps:

Step 1 Expand **All VSANs** and then select **IVR** in the Logical Domains pane.

Step 2 You see the IVR configuration in the Information pane shown in Figure 25-5.

Figure 25-5 IVR Routing Configuration Control Tab

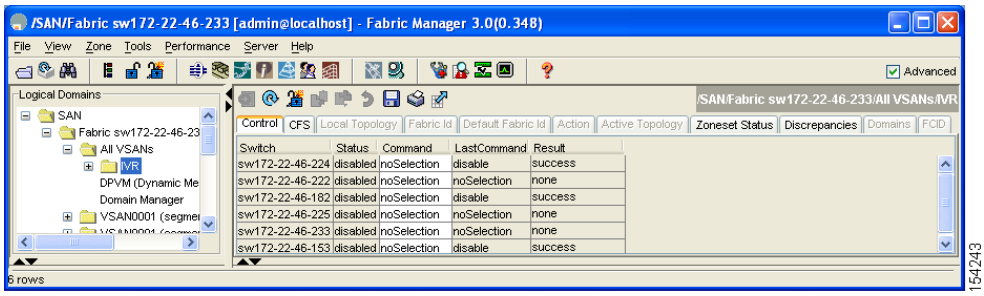

Step 3 Click the **Fabric ID** tab to display the existing AFIDs. (If this tab is greyed out, click the CFS tab first.)

Step 4 Click the **Create Row** icon to create an AFID. You see the AFID dialog box.

Step 5 Check the check box next to each switch involved in IVR that you want to use this default AFID.

Step 6 Provide a name for each SwitchWWN and set the Fabric ID.

Step 7 Enter a comma-separated list of VSAN IDs in the VSAN List text box.

Step 8 Click **Create** to create this entry or click **Cancel** to discard all changes.

Step 9 Repeat Step 1 through Step 7 for all switches and AFIDs you want to configure in your IVR topology.

About IVR Without IVR NAT or Auto Topology

Before configuring an IVR SAN fabric without IVR in NAT mode or IVR topology in auto mode, consider the following guidelines:

- Configure unique domain IDs across all VSANs and switches participating in IVR operations if you are not using IVR NAT. The following switches participate in IVR operations:
  - All edge switches in the edge VSANs (source and destination)
  - All switches in transit VSANs
- Configure IVR only in the relevant border switches.
- Acquire a mandatory Enterprise License Package or SAN-EXTENSION license package if you have Cisco MDS SAN-OS Release2.1(1a) or later and one active IPS card for this feature.

Tip If you change any FSPF link cost, ensure that the FSPF path distance (that is, the sum of the link costs on the path) of any IVR path is less than 30,000.

Note IVR-enabled VSANs can be configured when the interop mode is enabled (any interop mode) or disabled (no interop mode).

Domain ID Guidelines

Domain IDs must be unique across inter-connected VSANs when not using IVR NAT. To ensure unique domain IDs across inter-connected VSANs, consider these guidelines:

- Minimize the number of switches that require a domain ID assignment. This ensures minimum traffic disruption.
- Minimize the coordination between interconnected VSANs when configuring the SAN for the first time as well as when you add each new switch.

You can configure domain IDs using one of two options:

- Configure the allowed-domains list so that the domains in different VSANs are non-overlapping on all participating switches and VSANs.
- Configure static, non-overlapping domains for each participating switch and VSAN.

Note In a configuration involving IVR without NAT, if one VSAN in the IVR topology is configured with static domain IDs, then the other VSANs (edge or transit) in the topology should has to be configured with static domain IDs.

Transit VSAN Guidelines

Before configuring transit VSANS, consider the following guidelines:

- Besides defining the IVR zone membership, you can choose to specify a set of transit VSANs to provide connectivity between two edge VSANs:
  - If two edge VSANs in an IVR zone overlap, then a transit VSAN is not required (though, not prohibited) to provide connectivity.
  - If two edge VSANs in an IVR zone do not overlap, you may need one or more transit VSANs to provide connectivity. Two edge VSANs in an IVR zone will not overlap if IVR is not enabled on a switch that is a member of both the source and destination edge VSANs.
- Traffic between the edge VSANs only traverses through the shortest IVR path.
- Transit VSAN information is common to all IVR zone sets. Sometimes, a transit VSAN can also act as an edge VSAN in another IVR zone.

Border Switch Guidelines

Before configuring border switches, consider the following guidelines:

- Border switches require Cisco MDS SAN-OS Release 1.3(1) or later.
- A border switch must be a member of two or more VSANs.

- A border switch that facilities IVR communications must be IVR enabled.
- IVR can (optionally) be enabled on additional border switches to provide redundant paths between active IVR zone members.
- The VSAN topology configuration must be updated before a border switch is added or removed.

Configuring IVR Without NAT

To enable IVR without IVR in NAT mode from Fabric Manager, follow these steps:

Step 1 Expand **All VSANs** and then select **IVR** in the Logical Domains pane.

You see the IVR configuration in the Information pane shown in Figure 25-6.

Figure 25-6 IVR Routing Configuration Control Tab

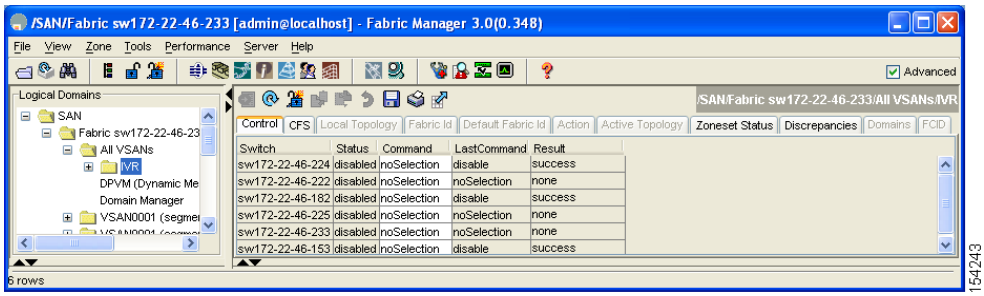

Step 2 Click the **CFS** tab if CFS is enabled for this feature in the fabric.

Step 3 Select **enable** from the Enable/Admin column for the primary switch.

Step 4 Click **Apply Changes** from the Information pane to distribute this change to all switches in the fabric, or click **Undo Changes** to cancel any changes you made.

Step 5 If CFS is not enabled, select the **Control** tab if it is not already displayed to enable IVR individually for each switch.

Step 6 Click **enable** from the command drop-down menu for each switch you want to enable IVR.

Step 7 Click **Apply Changes** from the Information pane to enable IVR on the switches, or click **Undo Changes** to cancel any changes you made.

Step 8 Click **CFS > Config Changes > Action** and choose **commit**.

Step 9 Click **Apply Changes** from the Information pane to distribute IVR on the switches.

Manually Creating the IVR Topology

You must create the IVR topology in every IVR-enabled switch in the fabric if you have not configured IVR topology in auto mode. You can have up to 128 VSANs in an IVR topology. Specify the IVR topology using the following information:

- The switch WWNs of the IVR-enabled switches.
- A minimum of two VSANs to which the IVR-enabled switch belongs.

- The AFID, which distinguishes two VSANs that are logically and physically separate, but have the same VSAN number (segmented VSANs). Cisco MDS SAN-OS Release 2.1(1a) supports up to 64 AFIDs.

Note Two VSANs with the same VSAN number but different AFIDs are counted as two VSANs out of the total 128 VSANs allowed in the fabric.

Note The use of a single AFID does not allow for segmented VSANs in an inter-VSAN routing topology.

To create the IVR topology using Fabric Manager, follow these steps:

Step 1 Expand **All VSANs** and then select **IVR** in the Logical Domains pane.

You see the IVR configuration in the Information pane shown in Figure 25-7.

Figure 25-7 IVR Routing Configuration Control Tab

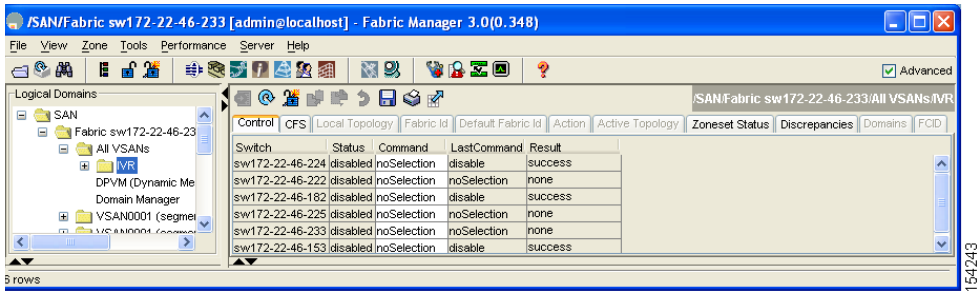

Step 2 Click the **Local Topology** tab to display the existing IVR topology.

Step 3 Click the **Create Row** icon from the Information pane to create rows in the IVR topology.

You see the local topology create dialog box.

Step 4 Select the switch, switch WWN, and a comma-separated list of VSAN IDs for this topology.

Step 5 Click **Create** to create this new row, or click **Close** to cancel all changes.

Step 6 Click **Apply Changes** from the Information pane to create the IVR topology, or click **Undo Changes** to cancel any changes you made.

Note Repeat this configuration in all IVR-enabled switches or distribute using CFS.

Tip Transit VSANs are deduced based on your configuration The IVR feature does not have an explicit transit-VSAN configuration.

Activating an IVR Topology

After creating the IVR topology, you must activate it.

To activate the IVR topology using Fabric Manager, follow these steps:

Step 1 Expand **All VSANs and then select IVR** in the Logical Domains pane.

You see the IVR configuration in the Information pane shown in Figure 25-8.

Figure 25-8 IVR Routing Configuration Control Tab

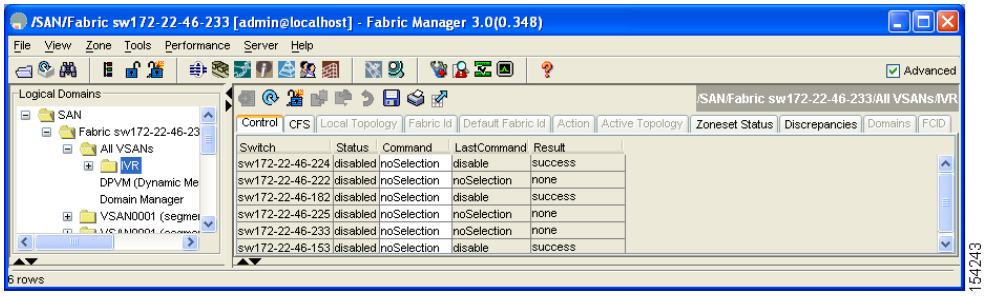

Step 2 Click the **Action** tab to display the existing IVR topology.

Step 3 Check the **Activate Local** check box.

Step 4 Click **Apply Changes** from the Information pane to activate the IVR topology, or click **Undo Changes** to cancel any changes you made.

Caution Active IVR topologies cannot be deactivated.

Clearing the IVR Topology

You can only clear manually created IVR VSAN topology entries from the config database.

To clear the IVR topology using Fabric Manager, follow these steps:

Step 1 Expand **All VSANs a**nd then select **IVR** in the Logical Domains pane.

You see the IVR configuration in the Information pane shown in Figure 25-9.

Figure 25-9 IVR Routing Configuration Control Tab

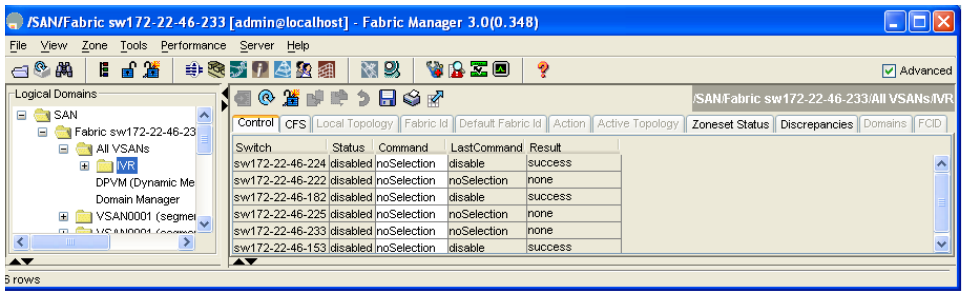

Step 2 Click the **Control** tab if it is not already displayed.

Step 3 Highlight the rows you want to delete from the IVR topology.

Step 4 Click the **Delete Row** icon from the Information pane to delete these rows from the IVR topology.

Step 5 Click **Apply Changes** from the Information pane to delete the IVR topology, or click **Undo Changes** to cancel the changes you made.

Migrating from IVR Auto Topology Mode to Manual Mode

If you want to migrate the active IVR VSAN topology database from automatic mode to user-configured mode, first copy the active IVR VSAN topology database to the user-configured IVR VSAN topology database before switching modes.

To migrate from automatic mode to manual mode using Fabric Manager, follow these steps:

Step 1 Expand **All VSANs** and then select **IVR** in the Logical Domains pane.

You see the IVR configuration in the Information pane.

Figure 25-10 IVR Routing Configuration Control Tab

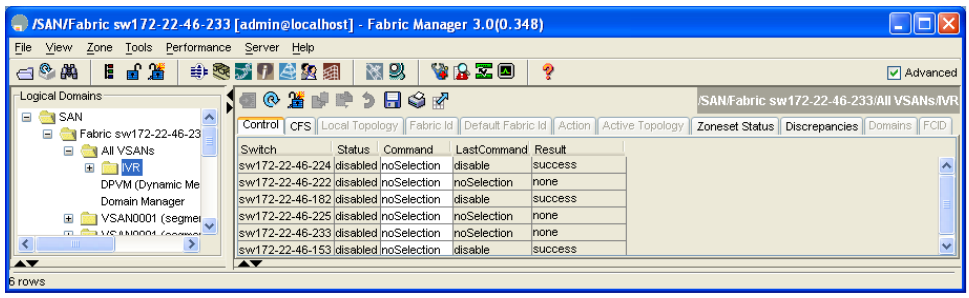

Step 2 Click the **CFS** tab to activate the rest of the tabs.

Step 3 Click the **Action** tab.

Step 4 Highlight the switch on which you want to disable auto topology mode.

Step 5 Uncheck the **Auto Discover Topology** check box.

Step 6 Click **Apply Changes** from the Information pane, or click **Undo Changes** to cancel any changes you made.

About IVR Virtual Domains

In a remote VSAN, the IVR application does not automatically add the virtual domain to the assigned domains list. Some switches (for example, the Cisco SN5428) do not query the remote name server until the remote domain appears in the assigned domains list in the fabric. In such cases, add the IVR virtual domains in a specific VSAN(s) to the assigned domains list in that VSAN. When adding IVR domains, all IVR virtual domains that are currently present in the fabric (and any virtual domain that is created in the future) will appear in the assigned domain list for that VSAN.

Tip Be sure to add IVR virtual domains if Cisco SN5428 or MDS 9020 switches exist in the VSAN.

When you enable the IVR virtual domains, links may fail to come up due to overlapping virtual domain identifiers. If so, temporarily withdraw the overlapping virtual domain from that VSAN.

Note Withdrawing an overlapping virtual domain from an IVR VSAN disrupts IVR traffic to and from that domain.

Configuring IVR Virtual Domains

To add IVR virtual domains using Fabric Manager, follow these steps:

Step 1 Expand **All VSANs** and then select **IVR** in the Logical Domains pane.

You see the IVR configuration in the Information pane shown in Figure 25-11.

Figure 25-11 IVR Routing Configuration Control Tab

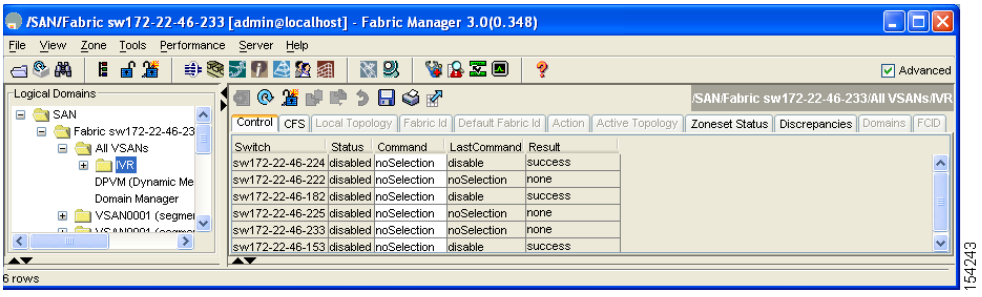

Step 2 Click the **CFS** tab to activate the rest of the tabs.

Step 3 Click the **Action** tab to display the existing IVR topology.

Step 4 Enter a comma-separated list of VSAN IDs in the Create Virtual Domain for VSANs column. These are the VSANs that will add the IVR virtual domains to the assigned domains list.

Step 5 Click **Apply Changes** from the Information pane to activate the IVR topology, or click **Undo Changes** to cancel any changes you made.

About Persistent FC IDs for IVR

You can configure persistent FC IDs for IVR. FC ID persistence improves IVR management by providing the following features:

- Allows you to control and assign a specific virtual domain to use for a native VSAN.
- Allows you to control and assign a specific virtual FC ID to use for a device.

The benefits of persistent FC IDs for IVR are as follows:

- Host devices always see the same FC ID for targets.
- It helps you plan your SAN layout better by assigning virtual domains for IVR to use.
- It can make SAN monitoring and management easier. When you see the same domain or FC ID consistently assigned, you can readily determine the native VSAN or device to which it refers.

You can configure two types of database entries for IVR FC ID persistence:

- Virtual domain entries—Contain the virtual domain that should be used to represent a native VSAN in a specific VSAN (current VSAN). These entries contain the following information:
  - Native AFID
  - Native VSAN
  - Current AFID
  - Current VSAN
  - Virtual domain to be used for the native AFID and VSAN in current AFID and VSAN
- Virtual FC ID entries—Contain the virtual FC ID that should be used to represent a device in a specific VSAN (current VSAN). These entries contain the following information:
  - Port WWN
  - Current AFID
  - Current VSAN
  - Virtual FC ID to be used to represent a device for the given pWWN in the current AFID and VSAN

Note In an IVR NAT configuration, if one VSAN in the IVR topology is configured with static domain IDs, then the IVR domains that can be exported to that VSAN must also be assigned static domains.

Configuring Persistent FC IDs for IVR

To configure persistent FC IDs for IVR using Fabric Manager, follow these steps:

Step 1 Expand **All VSANs** and then select **IVR** in the Logical Domains pane.

You see the IVR configuration in the Information pane shown in Figure 25-12.

Figure 25-12 IVR Routing Configuration Control Tab

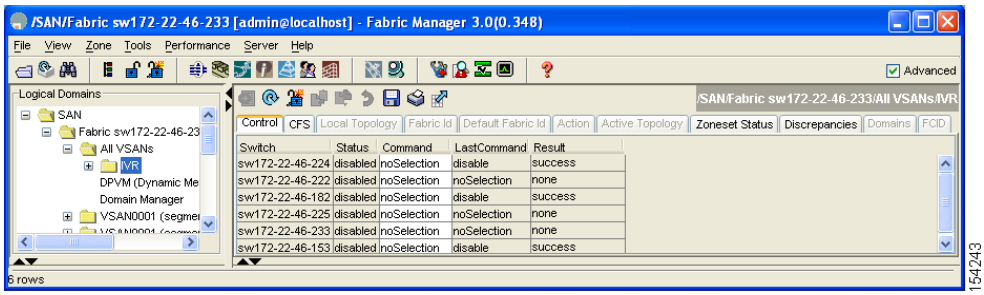

Step 2 Click the **CFS** tab to activate the rest of the tabs.

Step 3 Click the **FCID** tab.

Step 4 Click **Create Row** to create an FC ID.

Step 5 Select the switch for which you are configuring the virtual FC ID to be used to represent a device in a specific VSAN (current VSAN).

Step 6 Complete the Current Fabric ID field for the fcdomain database.

Step 7 Complete the Current VSAN ID field for the fcdomain database.

Step 8 Complete the pWWN field.

Step 9 Click the drop-down menu to select the FC ID to map to the pWWN you selected.

Step 10 Click **Create**.

Configuring IVR Logging Levels

To configure the severity level for logging messages from the IVR feature using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Events** and then select **Syslog** from the Physical Attributes pane.

You see the system messages configuration in the Information pane shown in Figure 25-13.

Figure 25-13 Syslog System Message Configuration

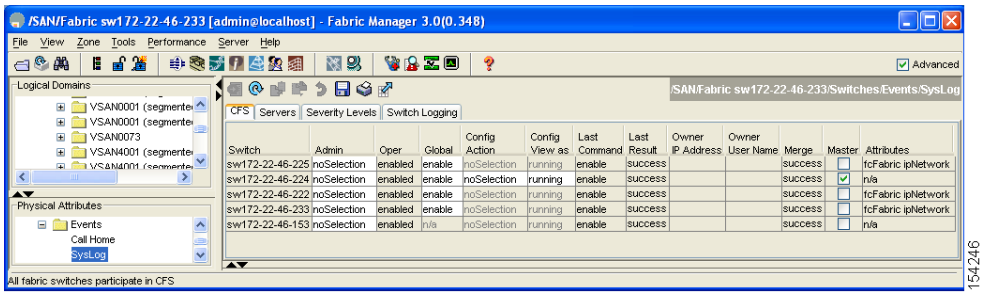

Step 2 Click the **Severity Levels** tab to see the list shown in Figure 25-14.

Figure 25-14 Severity Levels of Syslog Configured Messages

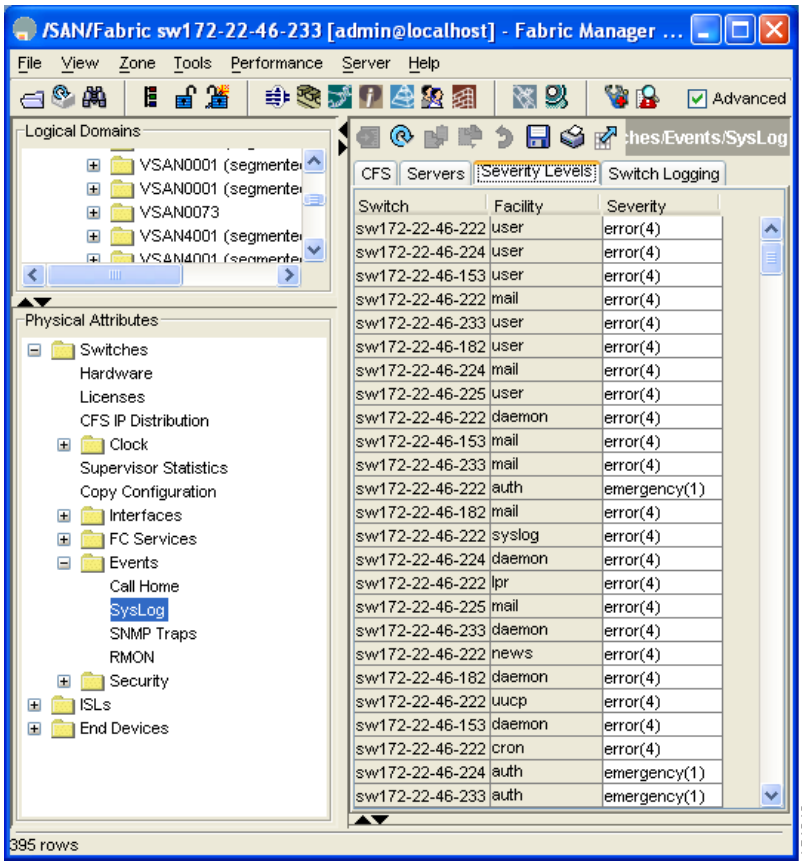

Step 3 Click the **Facility** column header to sort the table by facility name.

Step 4 Select the severity level at which the IVR logs system messages from the Severity drop-down menu shown in Figure 25-15.

Figure 25-15 Syslog Severity Drop-down Menu

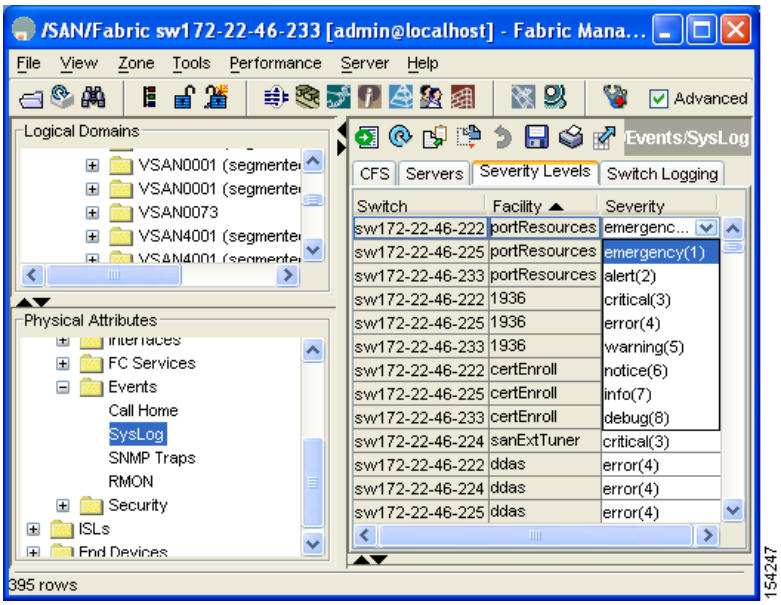

Tip Setting the severity to **warning** (option 5 in Figure 25-15) means that all IVR messages at the warning level or above will be logged to Fabric Manager.

Step 5 Click **Apply Changes** to save these changes locally.

Step 6 Click the **CFS** tab and select **commit** in the Config Action drop-down menu shown in Figure 25-16.

Figure 25-16 Config Action Drop-down Menu

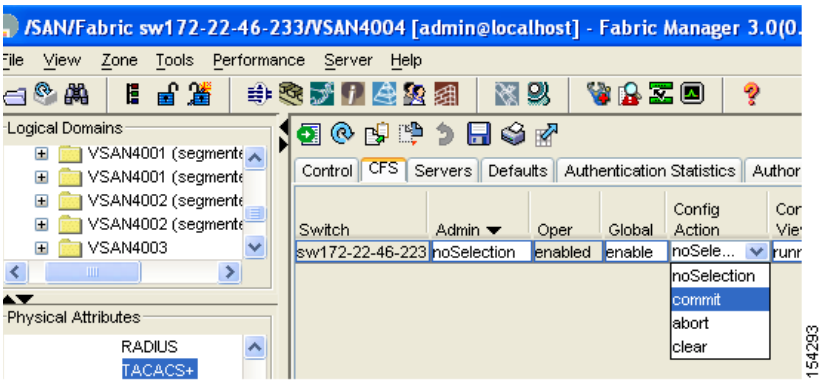

Step 7 Click **Apply Changes** to distribute these changes through the fabric.

IVR Zones and IVR Zone Sets

As part of the IVR configuration, you need to configure one or more IVR zone sets to enable cross-VSAN communication. To achieve this result, you must specify each IVR zone as a set of (pWWN, VSAN) entries. Like zones, several IVR zone sets can be configured to belong to an IVR zone. You can define several IVR zone sets and activate only one of the defined IVR zone sets.

Note The same IVR zone set must be activated on *all* of the IVR-enabled switches.

This section describes configuring IVR zones and IVR zone sets and includes the following topics:

About IVR Zones

Table 25-2 identifies the key differences between IVR zones and zones.

Table 25-2 Key Differences Between IVR Zones and Zones

IVR Zones	Zones
IVR zone membership is specified using the VSAN and pWWN combination.	Zone membership is specified using pWWN, fabric WWN, sWWN, or the AFID.
Default zone policy is always deny (not configurable).	Default zone policy is deny (configurable).

Automatic IVR Zone Creation

Figure 25-17 depicts an IVR zone consisting of four members. To allow pwwn1 to communicate with pwwn2, they must be in the same zone in VSAN 1, as well as in VSAN 2. If they are not in the same zone, then the hard-zoning ACL entries will prohibit pwwn1 from communicating with pwwn2.

A zone corresponding to each active IVR zone is automatically created in each edge VSAN specified in the active IVR zone. All pWWNs in the IVR zone are members of these zones in each VSAN.

Figure 25-17 Creating Zones Upon IVR Zone Activation

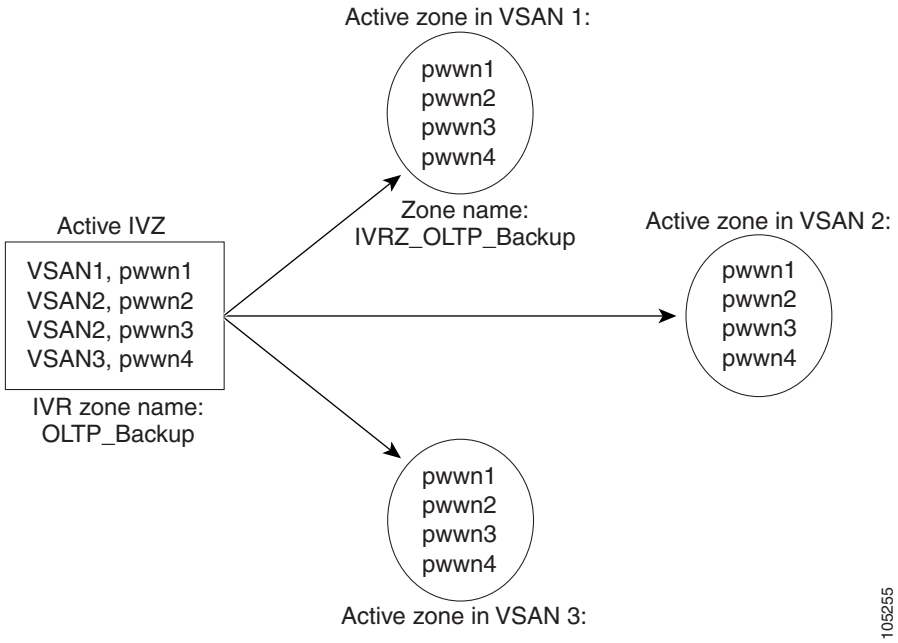

The zones are created automatically by the IVR process when an IVR zone set is activated. They are not stored in a full zone set database and are lost when the switch reboots or when a new zone set is activated. The IVR feature monitors these events and adds the zones corresponding to the active IVR zone set configuration when a new zone set is activated. Like zone sets, IVR zone sets are also activated nondisruptively.

Note If pwwn1 and pwwn2 are in an IVR zone in the current as well as the new IVR zone set , then activation of the new IVR zone set does not cause any traffic disruption between them.

IVR zone and IVR zone set names are restricted to 64 alphanumeric characters.

Configuring IVR Zones and Zone Sets

To create IVR zones or zone sets using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*** and then select **VSAN*xx*** for the VSAN that you want to configure in the Logical Domains pane.

Step 2 Click **Zone** and then select **Edit Local Full Zone Database** from the Zone menu (see Figure 25-18).

Figure 25-18 Zone Menu

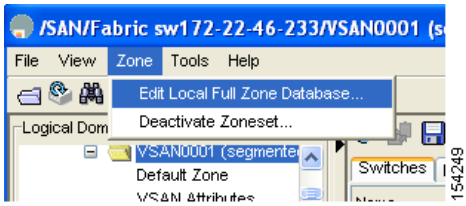

You see the Edit IVR Local Full Zone Database dialog box for the VSAN you selected (see Figure 25-19).

Figure 25-19 Edit Zone Database Dialog Box

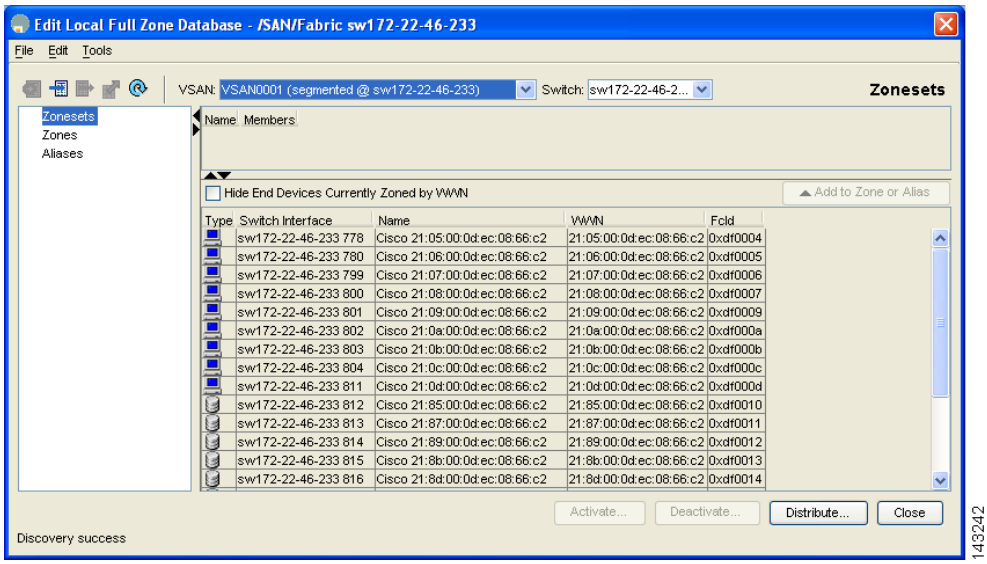

Step 3 Right-click either **Zones** or **Zonesets** in the left pane and select **Insert** to add a zone or zone set.

You see the dialog box in Figure 25-20.

Figure 25-20 Inserting a New Zone or Zone Set

Step 4 Provide a name then click **OK**.

You see the new folder in the left pane (Zoneset 1 in this example) shown in Figure 25-21.

Figure 25-21 New Zone Set Listed in Left Pane

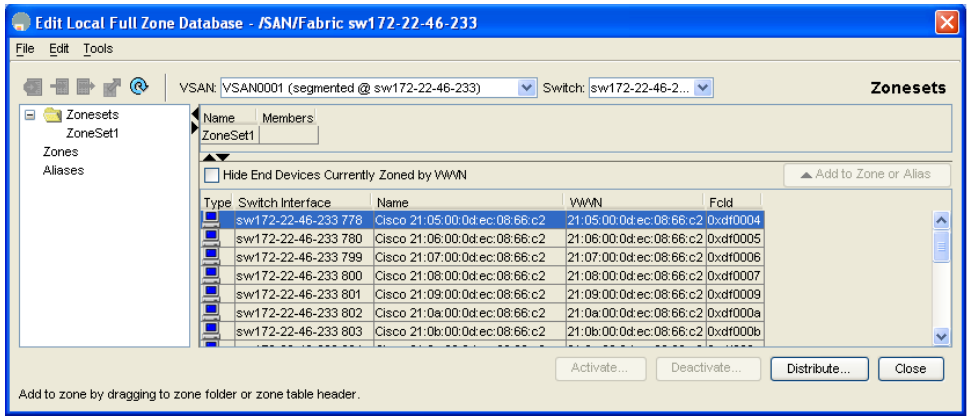

Step 5 If you added a zone set, drag existing zones into the new zone set in the left pane. If you added a zone, drag switches into the new zone in the left pane.

Step 6 If you added a zone set, select the new zone set and then click **Activate**.

You see the Save Configuration options shown in Figure 25-22.

Figure 25-22 Save Configuration Options for a New Zone Set

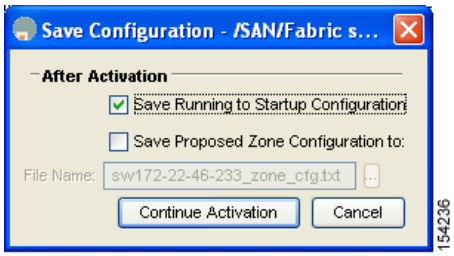

Step 7 Optionally check Save Running to Startup Configuration, or check Save Proposed Configuration and then supply a config name.

This permanently saves any changes made to the running or named configuration (not just zoning changes).

Step 8 Click **Continue Activation**.

Note Sometimes zone names beginning with prefix IVRZ and a zone set with name **nozoneset** appear in a logical view. The zones with prefix IVRZ are IVR zones that get appended to regular active zones. The prefix IVRZ is appended to active IVR zones by the system. Similarly the zone set with name **nozoneset** is an IVR active zone set created by the system if no active zone set is available for that VSAN and if the ivrZonesetActivateForce flag is enabled on the switch.

In the server.properties file, you can set the property zone.ignoreIVRZones to **true** or **false** to either hide or view IVR zones as part of regular active zones. See the "Fabric Manager Server Properties File" section on page 3-8 for more information on the server.properties file.

Note Do not create a zone with prefix the IVRZ or a zone set with name no zoneset. These names are used by the system for identifying IVR zones.

Step 9 Select the new zone or zone set from the list in the Information pane and then click **Distribute**.

About Zone Set Activation and the Force Activate Option

Once a zone set is created and populated, you activate the zone set (see Steps 6-8 in the previous procedure).

You can also use the force activate option to activate zone sets. Table 25-3 lists the various scenarios with and without the force activate option.

Table 25-3 IVR Scenarios with and without the Force Activate Option.

Case	Default Zone Policy	Active Zone Set before IVR Zone Activation	Force Activate Option Used?	IVR Zone Set Activation Status	Active IVR Zone Created?	Possible Traffic Disruption
1	Deny	No active zone set	No	Failure	No	No
2			Yes	Success	Yes	No
3[1]	**Deny**	**Active zone set present**	**No/Yes**	**Success**	**Yes**	**No**
4	Permit	No active zone set *or* Active zone set present	No	Failure	No	No
5			Yes	Success	Yes	Yes

1. We recommend that you use the Case 3 scenario.

Caution Using the force activate option may cause traffic disruption, even for devices that are not involved in IVR. For example, if your configuration does not have any active zone sets and the default zone policy is `permit`, then an IVR zone set activation will fail. However, IVR zone set activation will go through if the force activate option is used. Because zones are created in the edge VSANs corresponding to each IVR zone, traffic may be disrupted in edge VSANs where the default zone policy is `permit`.

Activating or Deactivating Zone Sets

To activate or deactivate an existing zone set using Fabric Manager, follow these steps:

Step 1 Click **Zone** and then select **Edit Local Full Zone Database** as shown in Figure 25-23.

Figure 25-23 Zone Menu

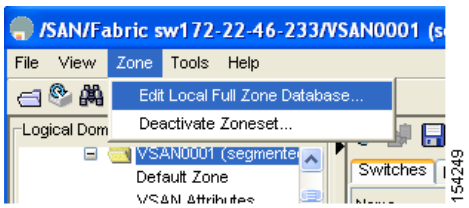

You see the Edit Local Full Zone Database dialog box in Figure 25-24.

Figure 25-24 Edit Zone Database Dialog Box

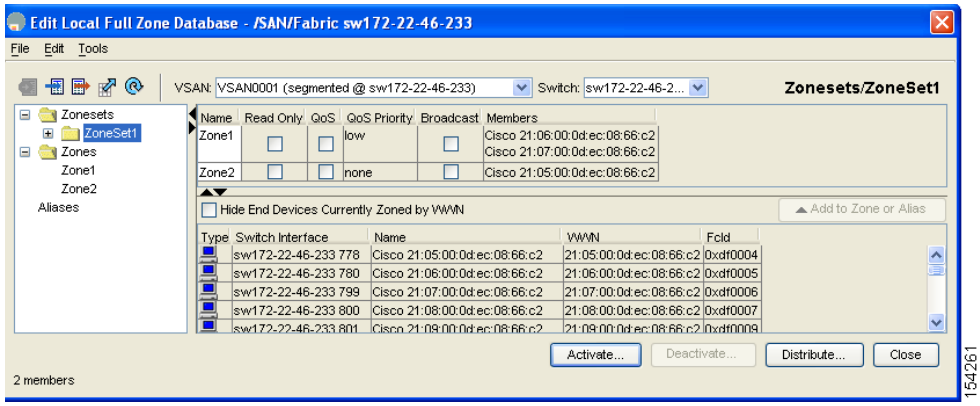

Step 2 Select a **Zoneset** folder and then click **Activate** to activate the zone set (shown in Figure 25-24) or click **Deactivate** to deactivate an activated zone set.

You see the Save Configuration dialog box shown in Figure 25-25.

Figure 25-25 Save Configuration Options for a New Zone Set

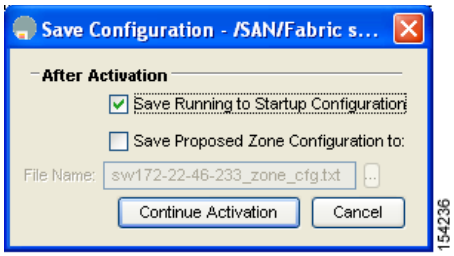

Step 3 Optionally, check one of the **Save Running to Configuration** check boxes to save these changes to the startup configuration (see Figure 25-25).

Step 4 Click **Continue Activation** to activate the zone set (see Figure 25-25) or **Yes** if you are deactivating the zone set.

Note The active zone set in Edit Zone is shown in bold if any change has been made to the full zoneset resulting in a difference between the active zoneset and full zoneset. Activating the zoneset, unbolds it.

Recovering an IVR Full Zone Database

You can recover an IVR zone database by copying the IVR full zone database from another switch.

To recover an IVR zone database using Fabric Manager, follow these steps:

Step 1 Click **Zone** and then select **Edit Local Full Zone Database** as shown in Figure 25-23.

You see the Edit Local Full Zone Database dialog box in Figure 25-24.

Step 2 Configure the Zone using the options for Read Only, QoS, QoS Priority, Broadcast, and Members.

Step 3 Click **Edit** and then select **Copy Full Zone Database** from the Edit menu (see Figure 25-26).

Figure 25-26 Edit Menu

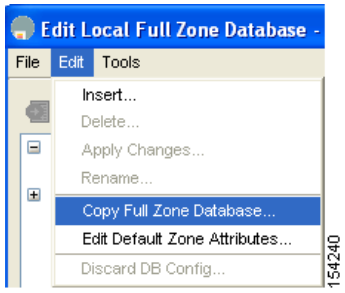

You see the Copy Full Zone Database dialog box shown in Figure 25-27.

Figure 25-27 Copy Full Zone Database Dialog Box

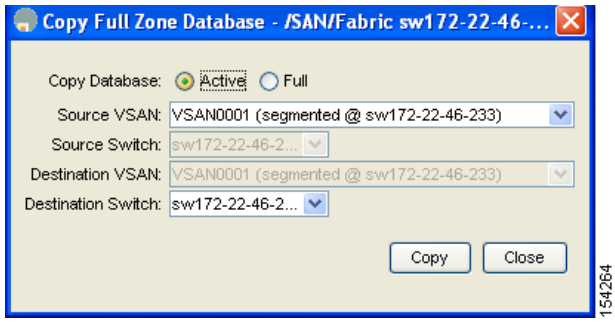

Step 4 Choose either **Active** or **Full**, depending on which type of IVR database you want to copy.

Step 5 Select the source switch from which to copy the information from the drop-down list.

Step 6 Select the destination switch from the drop-down list.

Step 7 Click **Copy** to copy the database, or click **Close** to close the dialog box without copying.

Recovering an IVR Full Topology

You can recover a topology by copying from the active zone database or the full zone database.

To recover a zone topology using Fabric Manager, follow these steps:

Step 1 Expand **SAN**, a fabric, and **ALL VSANS** in the Logical Domains pane.

Step 2 Right-click IVR then select **Copy Full Topology** (see Figure 25-28).

Figure 25-28 Copying a Full Topology

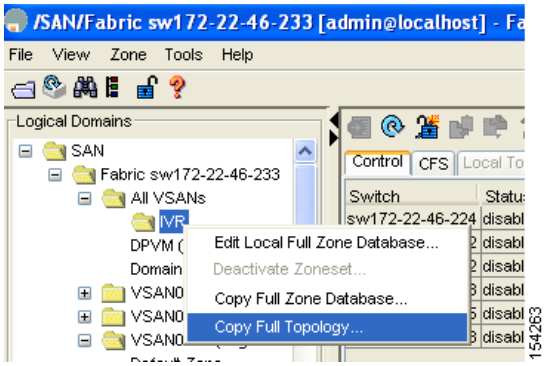

You see the Copy Full Topology dialog box.

Step 3 Select **Active** or **Full**, depending on which type of IVR database you want to copy from.

Step 4 Select the source switch from which to copy the information from the drop-down list.

Step 5 Select the destination switch from the drop-down list.

Step 6 Click **Copy** to copy the topology, or click **Close** to close the dialog box without copying.

Adding Members to IVR Zones

You can add members to existing IVR zones using the Edit Local Full Zone Database dialog box. LUN-zoning can optionally be used between members of active IVR zones.

To add members to an existing IVR zone and optionally configure LUN zoning using Fabric Manager, follow these steps:

Step 1 Click **Zone** and then select **Edit Local Full Zone Database**. You see the Edit Local Full Zone Database dialog box.

Step 2 Expand the **Zones** folder and select the zone to which you want to add a member.

Figure 25-29 Edit Zone Database Dialog Box

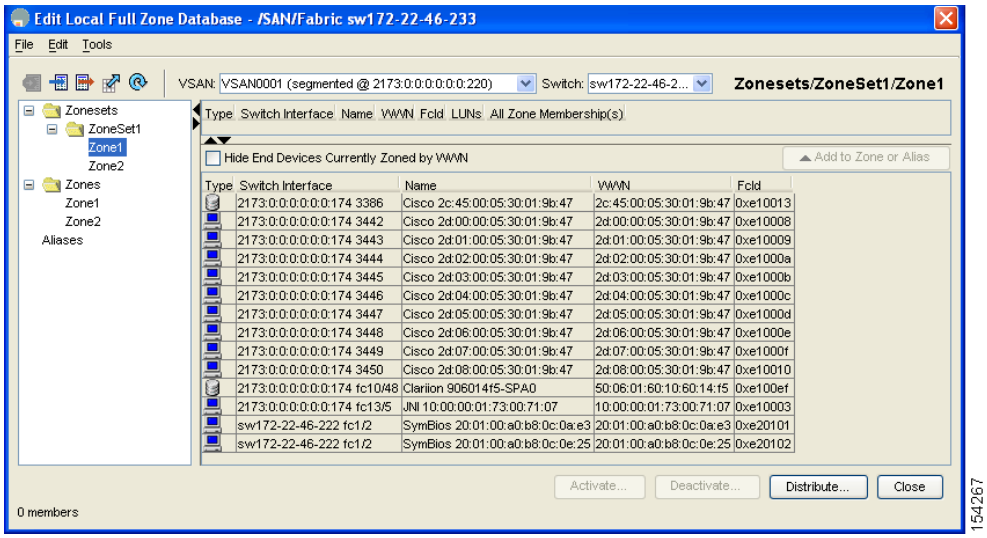

Step 3 Click the **Insert** icon to add a new member in this zone.

You see the zone membership dialog box. See Figure 25-30.

Figure 25-30 Add a Member to a Zone

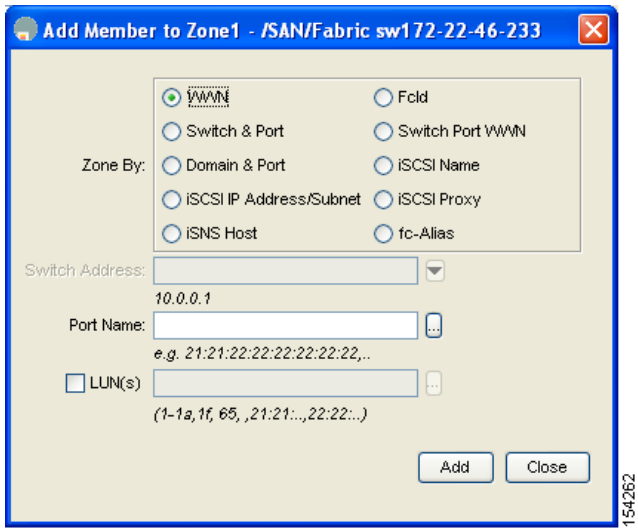

Step 4 Choose the zoning criteria from the options.

Step 5 Complete the Port Name field.

Step 6 Optionally, check the **LUNs** check box and set the LUNs you want this IVR zone to access on this end device.

Step 7 Click **Add** to add the member to the IVR zone with the optional LUN zoning attribute or click **Close** to discard all changes.

About LUNs in IVR Zoning

LUN zoning can be used between members of active IVR zones. Prior to Cisco MDS SAN-OS Release 2.1(1a), you can configure the service by creating and activating LUN zones between the desired IVR zone members in all relevant edge VSANs using the zoning interface. As of Cisco MDS SAN-OS Release 2.1(1a), IVR directly supports LUN zoning. For more details on the advantages of LUN zoning, see the "About LUN Zoning" section on page 26-38.

Configuring LUNs in IVR Zoning

As of Cisco MDS SAN-OS Release 2.1(1a), you can configure LUN zoning in an IVR zone set setup.

To configure LUNs in IVR zoning in Cisco MDS SAN-OS Release 2.1(1a) or later, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

About QoS in IVR Zones

As of Cisco MDS SAN-OS Release 2.1(1a), you can configure a QoS attribute for an IVR zone.

Configuring QoS for IVR Zones

To configure QoS for an IVR zone using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx*** and then select **VSAN*xx*** for the VSAN that you want to configure in the Logical Domains pane.

Step 2 Click **Zone** and then select **Edit Local Full Zone Database**.

You see the Edit IVR Local Full Zone Database dialog box for the VSAN you selected.

Step 3 Select **Zones** or a zone set.

Step 4 Check the **QoS** check box and set the QoS priority.

Step 5 Click **Activate** to make the changes or click **Close** to discard all changes.

Note If other QoS attributes are configured, the highest setting takes priority.

Renaming IVZs and IVZSs

You can rename IVZones and IVZone sets.

To rename an IVZ or IVRZset, using Fabric Manager, follow the steps below:

Step 1 Expand **Fabric*xx*** and then select **VSAN*xx*** for the VSAN that you want to configure in the Logical Domains pane.

Step 2 Choose **Zone** and then select **Edit Local Full Zone Database** from the Zone menu.

You see the Edit IVR Local Full Zone Database dialog box for the VSAN you selected.

Step 3 Right-click the zone set or zone for that VSAN.

Step 4 Click **Rename**.

Copying the Active IVZS

On the Cisco MDS Family switches, you cannot edit the active IVZS. However, you can copy the active IVZS to the full IVZS that you can edit.

To make a copy of the active IVZS using Fabric Manager, follow the steps below:

Step 1 Expand **Fabric*xx*** and then select **VSAN*xx*** for the VSAN that you want to configure in the Logical Domains pane.

Step 2 Choose **Zone** and then select **Edit Local Full Zone Database** from the Zone menu.

You see the Edit IVR Local Full Zone Database dialog box for the VSAN you selected.

Step 3 Right-click a zone set.

Step 4 Click **Copy**.

You see the Copy ZoneSet Dialog Box

Step 5 You are given the option to prepend or append the IVR Zoneset with the selected name.

Step 6 Click **OK** to proceed or click **Cancel**.

Clearing the IVR Zone Database

Clearing a zone set only erases the configured zone database, not the active zone database.

To clear the IVR zone database, refer to the *Cisco MDS 9000 CLI Configuration Guide*.

Configuring IVR Using Read-Only Zoning

Read-only zoning (with or without LUNs) can be used between members of active IVR zones. To configure this service, you must create and activate read-only zones between the desired IVR zone members in all relevant edge VSANs using the zoning interface.

Note Read-only zoning cannot be configured in an IVR zone set setup.

Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the active database. See the "CFS Merge Support" section on page 12-9 for detailed concepts.

- Be aware of the following conditions when merging two IVR fabrics:
  - The IVR configurations are merged even if two fabrics contain different configurations.
  - If dissimilar zones exist in two merged fabrics, the zone from each fabric is cloned in the distributed zone set with appropriate names (see Figure 25-31).

Figure 25-31 Fabric Merge Consequences

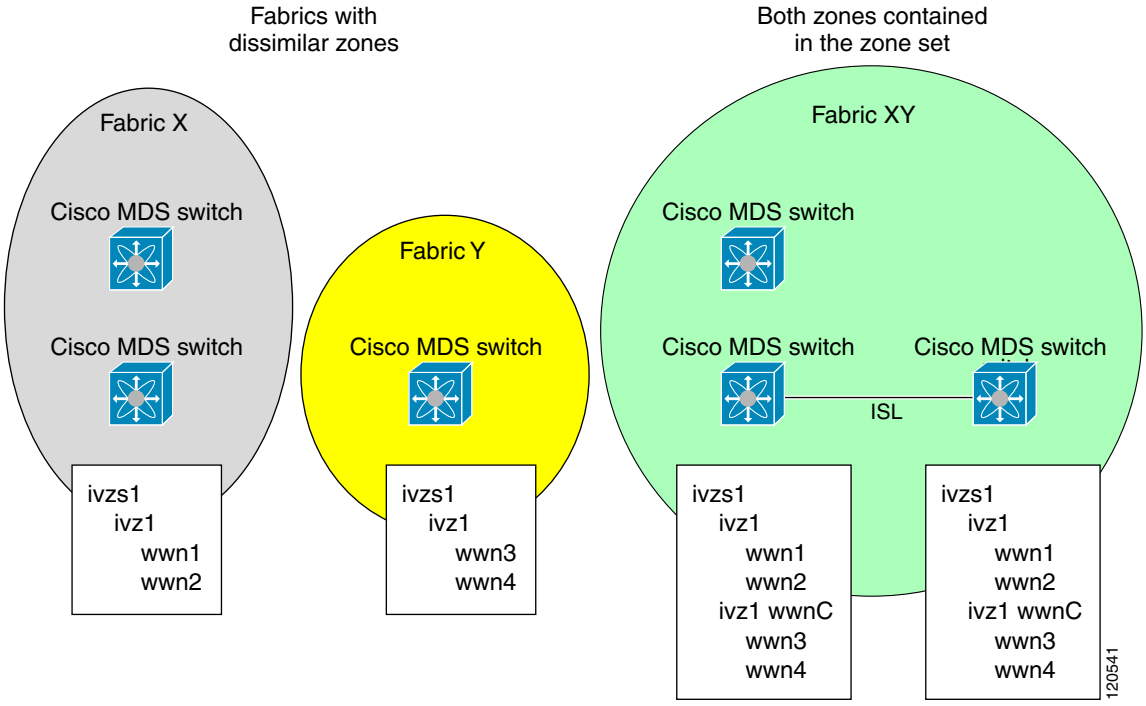

- You can configure different IVR configurations in different Cisco MDS switches.
- Be aware that the merge follows more liberal approach in order to avoid traffic disruption. After the merge, the configuration will be a union of the configurations that were present on the two switches involved in the merge.
  - The configurations are merged even if both fabrics have different configurations.
  - A union of zones and zone sets are used to get the merged zones and zone sets. If a dissimilar zone exists in two fabrics, the dissimilar zones are cloned into the zone set with appropriate names so both zones are present.
  - The merged topology contains a union of the topology entries for both fabrics.
  - The merge will fail if the merged database contains more topology entries than the allowed maximum.
  - The total number of VSANs across the two fabrics cannot exceed 64. As of Cisco MDS SAN-OS Release 2.1(1a), the total number of VSANs across the two fabrics cannot exceed 128.

Note VSANs with the same VSAN ID but different AFIDs are counted as two separate VSANs.

  - The total number of IVR-enabled switches across the two fabrics cannot exceed 128.
  - The total number of zone members across the two fabrics cannot exceed 2000. As of Cisco MDS SAN-OS Release 2.1(1a), the total number of zone members across the two fabrics cannot exceed 10,000. A zone member is counted twice if it exists in two zones.
  - The total number of zones across the two fabrics cannot exceed 200. As of Cisco MDS SAN-OS Release 2.1(1a), the total number of zones across the two fabrics cannot exceed 2000.
  - The total number or zone sets across the two fabrics cannot exceed 32.

Table 25-4 describes the results of a CFS merge of two IVR-enabled fabrics under different conditions.

Table 25-4 Results of Merging Two IVR-Enabled Fabrics

IVR Fabric 1	IVR Fabric 2	After Merge
NAT enabled	NAT disabled	Merge succeeds and NAT enabled
Auto mode on	Auto mode off	Merge succeeds and auto mode on
Conflicting AFID database		Merge fails
Conflicting IVR zone set database		Merge succeeds with new zones created to resolve conflicts
Combined configuration exceeds limits (such as maximum number of zones or VSANs)		Merge fails
Service group 1	Service group 2	Merge succeeds with service groups combined
User-configured VSAN topology configuration with conflicts		Merge fails
User-configured VSAN topology configuration without conflicts		Merge succeeds

Caution If you do not follow these conditions, the merge will fail. The next distribution will forcefully synchronize the databases and the activation states in the fabric.

Default Settings

Table 25-5 lists the default settings for IVR parameters.

Table 25-5 Default IVR Parameters

Parameters	Default
IVR feature	Disabled.
IVR VSANs	Not added to virtual domains.
IVR NAT	Disabled.
QoS for IVR zones	Low.
Configuration distribution	Disabled.

CHAPTER 26

Configuring and Managing Zones

Zoning enables you to set up access control between storage devices or user groups. If you have administrator privileges in your fabric, you can create zones to increase network security and to prevent data loss or corruption. Zoning is enforced by examining the source-destination ID field.

Advanced zoning capabilities specified in the FC-GS-4 and FC-SW-3 standards are provided. You can use either the existing basic zoning capabilities or the advanced, standards-compliant zoning capabilities.

This chapter includes the following sections:

Note Table 23-1 on page 23-4 lists the differences between zones and VSANs.

About Zoning

Zoning has the following features:

- A zone consists of multiple zone members.
  - Members in a zone can access each other; members in different zones cannot access each other.
  - If zoning is not activated, all devices are members of the default zone.
  - If zoning is activated, any device that is not in an active zone (a zone that is part of an active zone set) is a member of the default zone.
  - Zones can vary in size.
  - Devices can belong to more than one zone.
- A zone set consists of one or more zones.
  - A zone set can be activated or deactivated as a single entity across all switches in the fabric.
  - Only one zone set can be activated at any time.
  - A zone can be a member of more than one zone set.
- Zoning can be administered from any switch in the fabric.
  - When you activate a zone (from any switch), all switches in the fabric receive the active zone set. Additionally, full zone sets are distributed to all switches in the fabric, if this feature is enabled in the source switch.
  - If a new switch is added to an existing fabric, zone sets are acquired by the new switch.
- Zone changes can be configured nondisruptively. New zones and zone sets can be activated without interrupting traffic on unaffected ports or devices.
- Zone membership criteria is based on WWNs or FC IDs.
  - Port world wide name (pWWN)—Specifies the pWWN of an N port attached to the switch as a member of the zone.
  - Fabric pWWN—Specifies the WWN of the fabric port (switch port's WWN). This membership is also referred to as port-based zoning.
  - FC ID—Specifies the FC ID of an N port attached to the switch as a member of the zone.
  - Interface and switch WWN (sWWN)—Specifies the interface of a switch identified by the sWWN. This membership is also referred to as interface-based zoning.
  - Interface and domain ID—Specifies the interface of a switch identified by the domain ID.
  - Domain ID and port number—Specifies the domain ID of an MDS domain and additionally specifies a port belonging to a non-Cisco switch.
  - IP address—Specifies the IP address (and optionally the subnet mask) of an attached device.
- IPv6 address—The IPv6 address of an attached device in 128 bits in colon(:)-separated hexadecimal format.
- Default zone membership includes all ports or WWNs that do not have a specific membership association. Access between default zone members is controlled by the default zone policy.
- You can configure up to 8000 zones per VSAN and a maximum of 8000 zones for all VSANs on the switch.

Zoning Example

Figure 26-1 illustrates a zone set with two zones, zone 1 and zone 2, in a fabric. Zone 1 provides access from all three hosts (H1, H2, H3) to the data residing on storage systems S1 and S2. Zone 2 restricts the data on S3 to access only by H3. Note that H3 resides in both zones.

Figure 26-1 Fabric with Two Zones

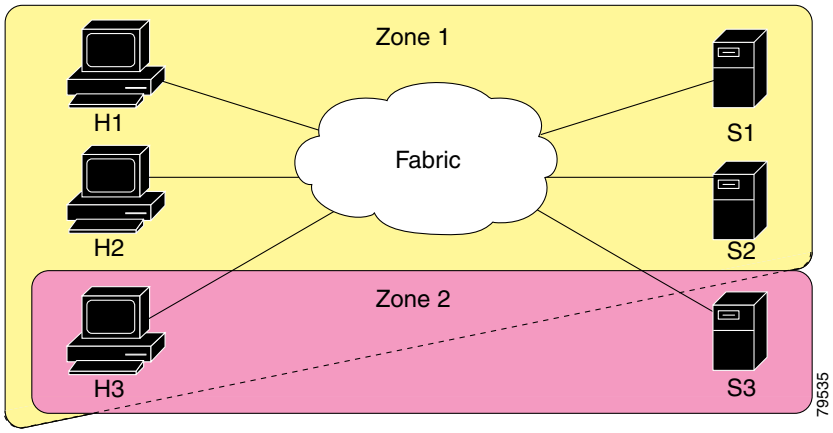

Of course, there are other ways to partition this fabric into zones. Figure 26-2 illustrates another possibility. Assume that there is a need to isolate storage system S2 for the purpose of testing new software. To achieve this, zone 3 is configured, which contains only host H2 and storage S2. You can restrict access to just H2 and S2 in zone 3, and to H1 and S1 in zone 1.

Figure 26-2 Fabric with Three Zones

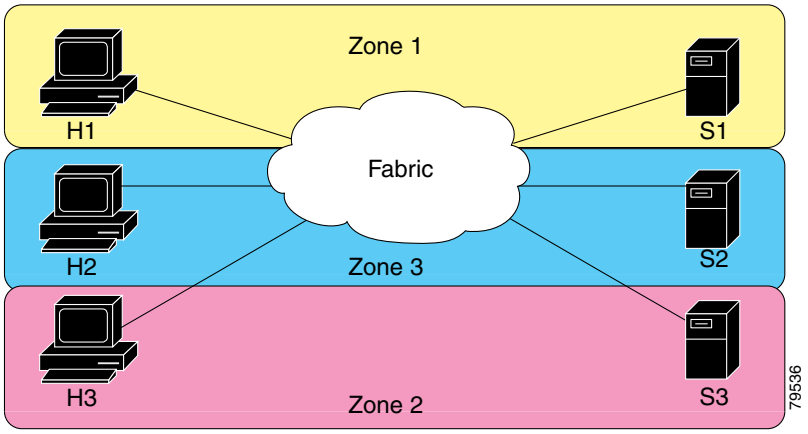

Zone Implementation

All switches in the Cisco MDS 9000 Family automatically support the following basic zone features (no additional configuration is required):

- Zones are contained in a VSAN.
- Hard zoning cannot be disabled.
- Name server queries are soft-zoned.
- Only active zone sets are distributed.
- Unzoned devices cannot access each other.
- A zone or zone set with the same name can exist in each VSAN.
- Each VSAN has a full database and an active database.
- Active zone sets cannot be changed, without activating a full zone database.
- Active zone sets are preserved across switch reboots.
- Changes to the full database must be explicitly saved.
- Zone reactivation (a zone set is active and you activate another zone set) does not disrupt existing traffic.

If required, you can additionally configure the following zone features:

- Propagate full zone sets to all switches on a per VSAN basis.
- Change the default policy for unzoned members.
- Interoperate with other vendors by configuring a VSAN in the interop mode. You can also configure one VSAN in the interop mode and another VSAN in the basic mode in the same switch without disrupting each other
- Bring E ports out of isolation.

Zone Configuration

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

About Zone Configuration

A zone can be configured using one of the following identifiers to assign members:

- pWWN—The WWN of the N or NL port is in hex format (for example, 10:00:00:23:45:67:89:ab).

Caution You must only configure pWWN-type zoning on all MDS switches running Cisco SAN-OS if there is a Cisco MDS 9020 switch running FabricWare in the same fabric.

- Fabric port WWN—The WWN of the fabric port name is in hex format (for example, 10:00:00:23:45:67:89:ab).
- FC ID—The N port ID is in 0xhhhhhh format (for example, 0xce00d1).
- FC alias—The alias name is in alphabetic characters (for example, Payroll) and denotes a port ID or WWN. The alias can also include multiple members.
- Domain ID—The domain ID is an integer from 1 to 239. A mandatory port number of a non-Cisco switch is required to complete this membership configuration.
- IP address—The IP address of an attached device in 32 bytes in dotted decimal format along with an optional subnet mask. If a mask is specified, any device within the subnet becomes a member of the specified zone.
- IPv4 address—The IPv4 address of an attached device in 32 bits in dotted decimal format along with an optional subnet mask. If a mask is specified, any device within the subnet becomes a member of the specified zone.
- IPv6 address—The IPv6 address of an attached device in 128 bits in colon(:)-separated hexadecimal format.
- Interface—Interface-based zoning is similar to port-based zoning because the switch interface is used to configure the zone. You can specify a switch interface as a zone member for both local and remote switches. To specify a remote switch, enter the remote switch WWN (sWWN) or the domain ID in the particular VSAN.

About the Edit Full Zone Database Tool

For version 2.0, there are interface changes to the Edit Full Zone Database screen, which is shown in Figure 26-3.

Figure 26-3 Edit Full Zone Database Screen

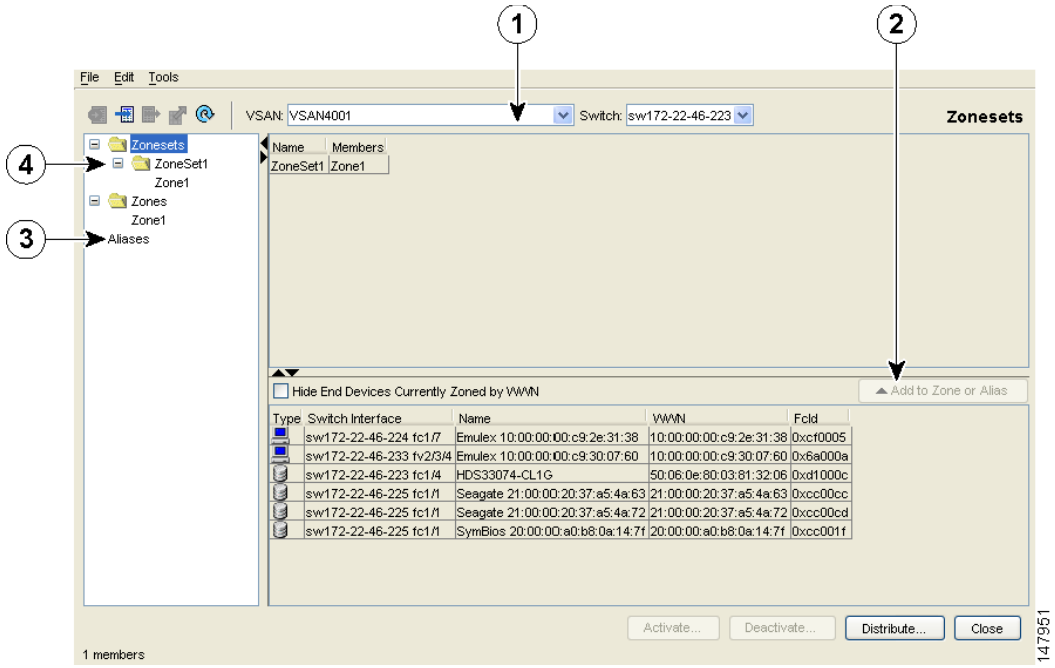

1	You can display information by VSAN by using the pull-down menu without having to get out of the screen, selecting a VSAN, and re-entering.	**3**	You can add zoning characteristics based on alias in different folders.
2	You can use the **Add to zone or alias** button to move devices up or down by alias or by zone.	**4**	You can triple-click to rename zone sets, zones, or aliases in the tree.

Tip Expand **Switches** from the Physical Attributes pane to retrieve the sWWN. If you do not provide a sWWN, the software automatically uses the local sWWN.

Note Interface-based zoning only works with Cisco MDS 9000 Family switches. Interface-based zoning does not work if interop mode is configured in that VSAN.

Configuring a Zone Using the Zone Configuration Tool

To configure a zone and assign a zone name using Fabric Manager, follow these steps:

Step 1 Click the **Zone** icon in the toolbar. (See Figure 26-4.)

Figure 26-4 Zone Icon

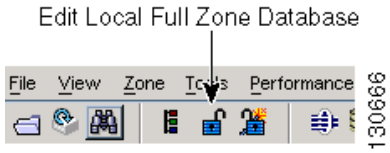

Step 2 Select the VSAN where you will configure zone sets, zones, or add members to a zone.

You see the Edit Local Full Zone Database dialog box shown in Figure 26-5.

Figure 26-5 Edit Local Full Zone Database

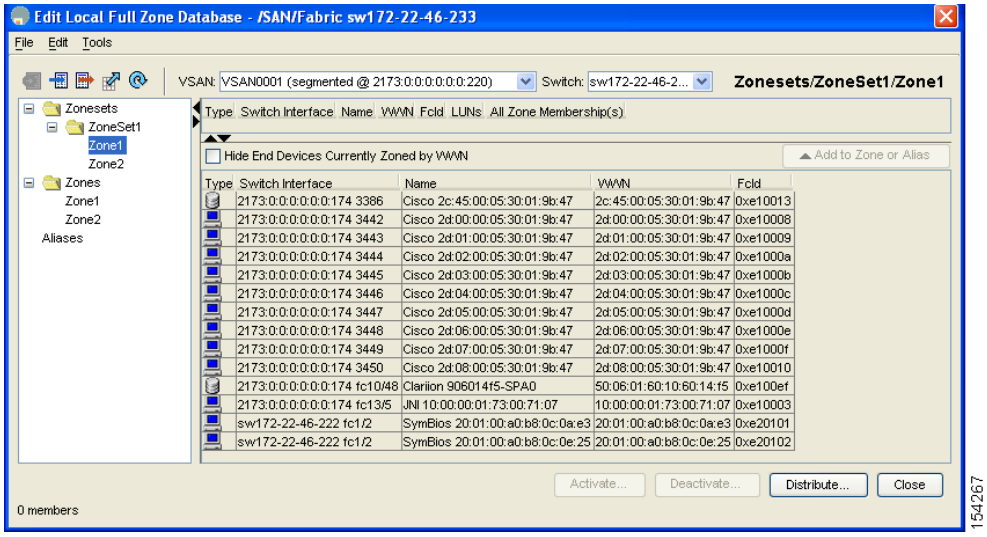

Step 3 Select **Zoneset** in the left pane and then click the **Create Row** icon to create a new zone set.

You see the Zoneset Name dialog box in Figure 26-6.

Figure 26-6 Zoneset Name Dialog Box

Step 4 Provide a zone set name and then click **OK**.

Step 5 Select **Zones** in the left pane and click the **Create Row** icon to make a new zone.

You see the Create Zone dialog box in Figure 26-7.

Figure 26-7 Create Zone Dialog Box

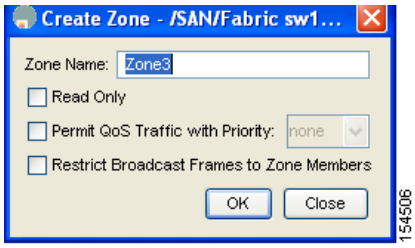

a. Optionally, check the **Read Only** check box if you want the zone to permit read and deny writes.

b. Optionally, check the **Permit QoS traffic with Priority** check box and set the QoS priority from the drop-down menu.

c. Optionally, check the **Restrict Broadcast frames to Zone Members** check box.

Step 6 Click **OK** to create the zone.

Step 7 Select **Aliases** in the left pane and then click **Create Row** to make a new device alias.

You see the Create Alias dialog box shown in Figure 26-8.

Figure 26-8 Create Alias Dialog Box

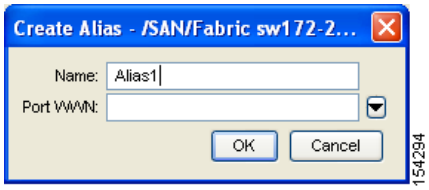

Step 8 Provide the two names and then click **OK**.

Adding Zone Members

Once you have created a zone, you can add members to the zone. You can add members using multiple port identification types. See the "Zone Configuration" section on page 26-4.

To add a member to a zone using Fabric Manager, follow these steps:

Step 1 Click the **Zone menu** and then select **Edit Local Full Zone Database** as shown in Figure 26-13, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the drop-down menu.

Figure 26-9 Click Zone and then select Edit Local Full Zone Database

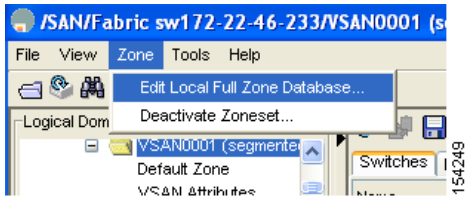

If you chose **Zone > Edit Local Full Zone Database**, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

Either way, you see the Edit Local Full Zone Database window for the VSAN you selected (see Figure 26-14).

Figure 26-10 Edit Local Full Zone Database

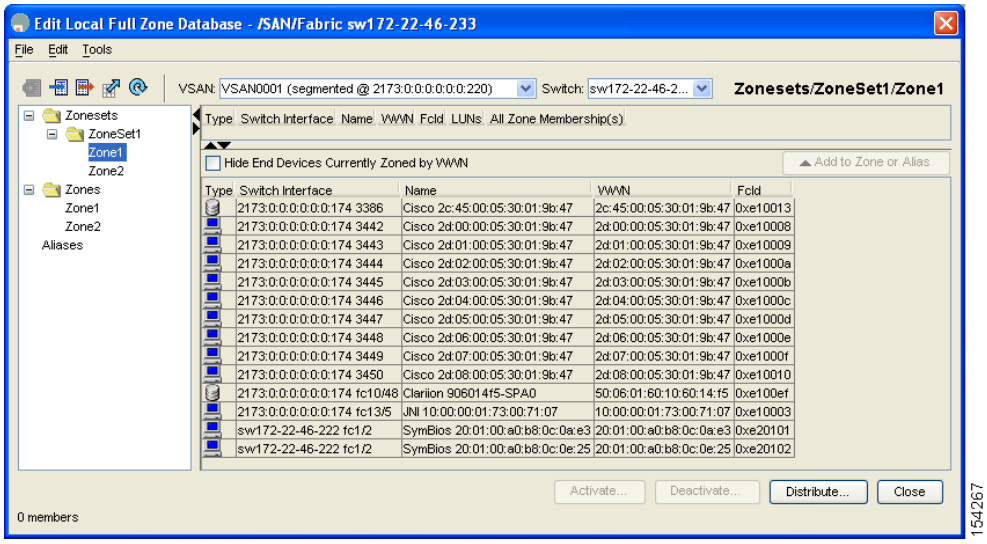

Step 2 Right-click the folder for a zone to which you want to add members, and then choose **Insert** from the drop-down menu (see Figure 26-11).

Figure 26-11 Right-click a Folder and Select Insert to Add a Zone Member

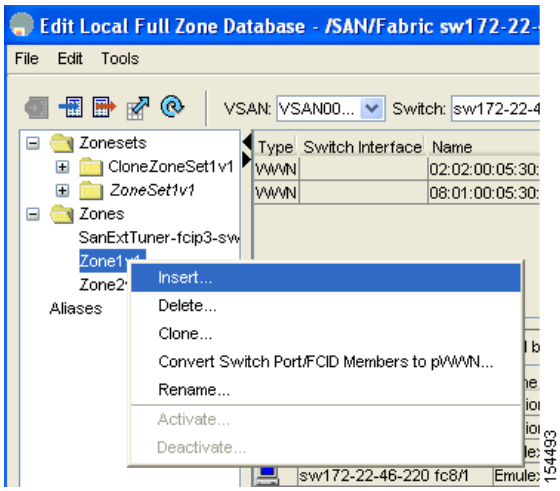

You see the Add Member to Zone dialog box shown in Figure 26-12.

Figure 26-12 Add Member to Zone

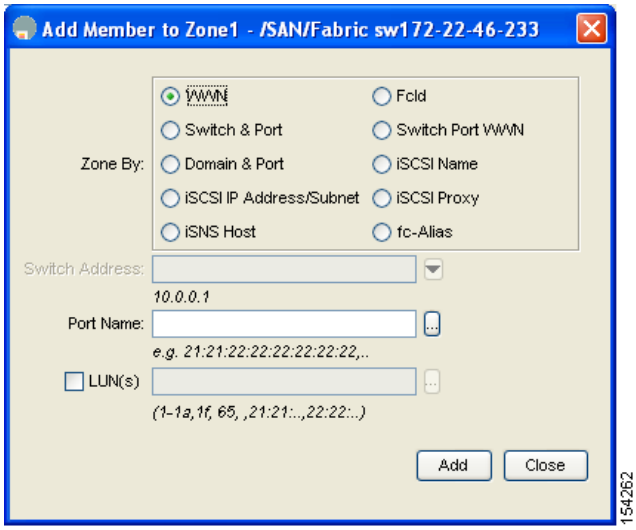

Step 3 Optionally, change the **Zone By** option (see Figure 26-12). Zoning is done by WWN by default.

Step 4 Select one of the port identifier options in the dialog box and click **Add** to add it to the zone (see Figure 26-12).

You see the member in the zone server database in the lower frame.

Optionally, check the **LUN** check box and click**...** to configure LUNS.

Step 5 Repeat these steps to add other members to the zone.

Note When configuring a zone member, you can specify that a single LUN has multiple IDs depending on the operating system. You can select from six different operating systems.

Displaying Zone Membership Information

To display zone membership information for members assigned to zones in Fabric Manager, follow these steps:

Step 1 Click **Zone** and then select **Edit Local Full Zone Database** from the Zone menu shown in Figure 26-13, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the drop-down menu.

Figure 26-13 Edit Local Full Zone Database

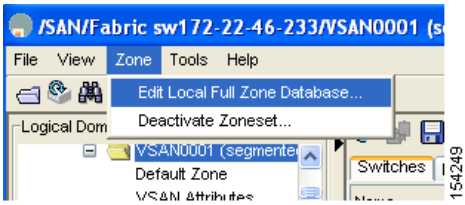

If you chose **Zone > Edit Local Full Zone Database**, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

Either way, you see the Edit Local Full Zone Database window for the VSAN you selected (see Figure 26-14).

Figure 26-14 Edit Local Full Zone Database Dialog Box

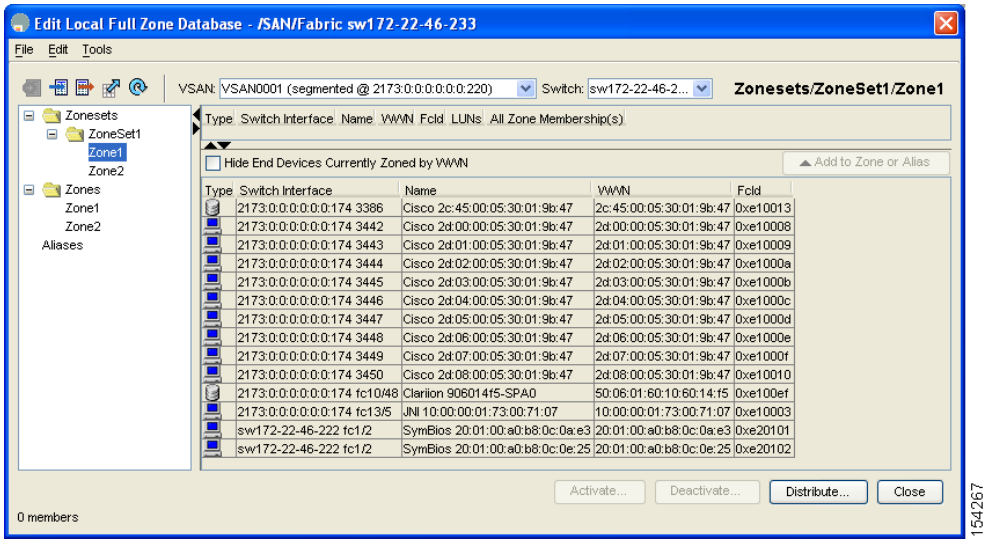

Step 2 Select a **Zones** folder. (Zone1 is selected in Figure 26-14.)

The right pane lists the members for each zone.

Note The default zone members are explicitly listed only when the default zone policy is configured as **permit**. When the default zone policy is configured as **deny**, the members of this zone are not shown. See the "Displaying Zone Information" section on page 26-42.

About Alias Creation

You can assign an alias name and configure an alias member using either the FC ID, fabric port WWN (fWWN), or pWWN values.

- pWWN—The WWN of the N or NL port in hex format (for example, 10:00:00:23:45:67:89:ab).
- fWWN—The WWN of the fabric port name in hex format (for example, 10:00:00:23:45:67:89:ab).
- FC ID—The N port ID in 0xhhhhhh format (for example, 0xce00d1).
- Domain ID—The domain ID is an integer from 1 to 239. A mandatory port number of a non-Cisco switch is required to complete this membership configuration.
- IPv4 address—The IPv4 address of an attached device in 32 bits in dotted decimal format along with an optional subnet mask. If a mask is specified, any device within the subnet becomes a member of the specified zone.
- IPv6 address—The IPv6 address of an attached device in 128 bits in colon(:)-separated hexadecimal format.
- Interface—Interface-based zoning is similar to port-based zoning because the switch interface is used to configure the zone. You can specify a switch interface as a zone member for both local and remote switches. To specify a remote switch, enter the remote switch WWN (sWWN) or the domain ID in the particular VSAN.

Tip The Cisco SAN-OS software supports a maximum of 2048 aliases per VSAN.

Creating Aliases

To create an alias using Fabric Manager, follow these steps:

Step 1 Click **Zone** and then select **Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the drop-down menu.

If you chose **Zone > Edit Local Full Zone Database** , then you see the Select VSAN dialog box . Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window for the VSAN you selected (see Figure 26-15).

Figure 26-15 Edit Local Full Zone Database

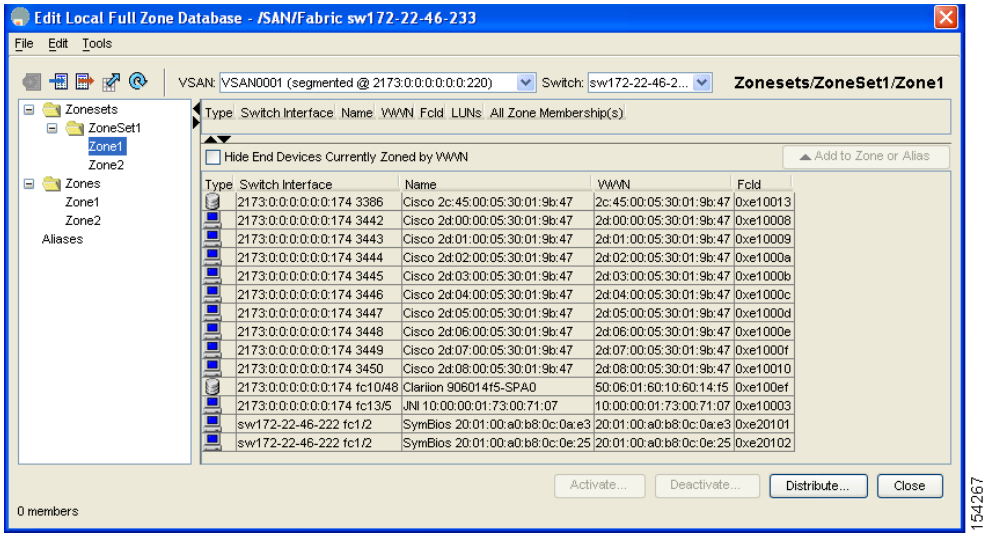

Step 2 Select the **Aliases** folder in the lower left pane.

The right pane lists the existing aliases.

Step 3 Click the **Insert** icon to create an alias.

You see the Create Alias dialog box shown in Figure 26-16.

Figure 26-16 Create Alias

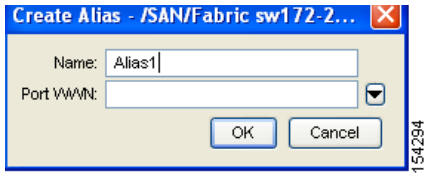

Step 4 Supply the pWWN and Alias Name (see Figure 26-16).

Step 5 Click **OK** to create the alias or click **Close** to close the dialog box (see Figure 26-16).

Adding Members to Aliases

To add a member to an alias using Fabric Manager, follow these steps:

Step 1 Click **Zone** and then select **Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the drop-down menu.

If you chose **Zone > Edit Local Full Zone Database** , then you see the Select VSAN dialog box . Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window for the VSAN you selected.

Step 2 Select the **Aliases** folder in the left pane (see Figure 26-17).

The right pane lists the existing aliases (see Figure 26-17).

Figure 26-17 Aliases

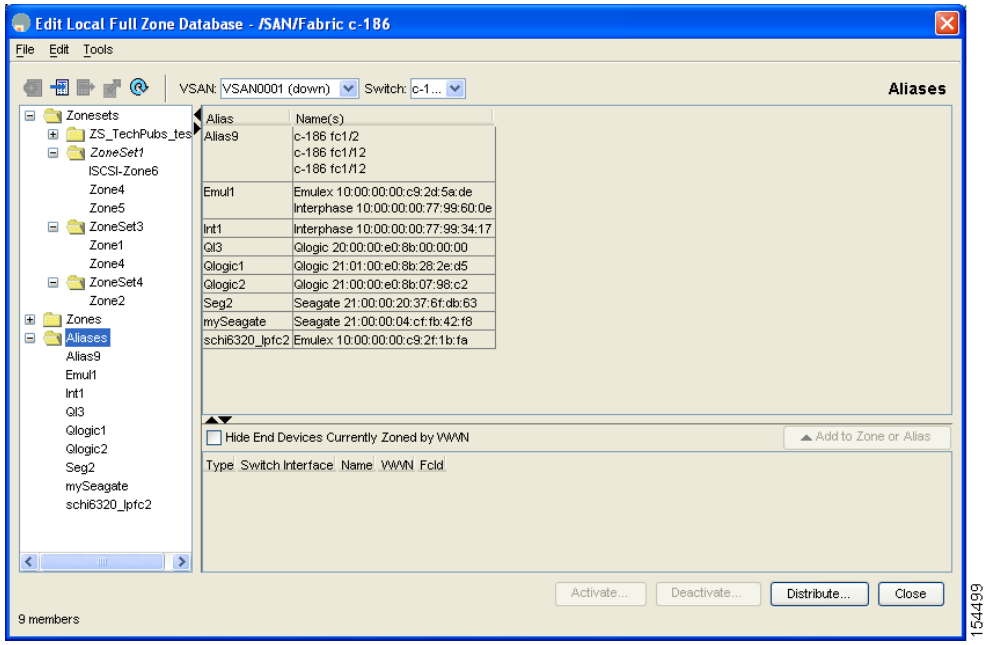

Step 3 Click the alias that you want to add a member (see Figure 26-17) to and then click the **Insert** icon to add a new member.

You see the Add Member to Alias dialog box Figure 26-18.

Figure 26-18 Add Member to Alias Dialog Box

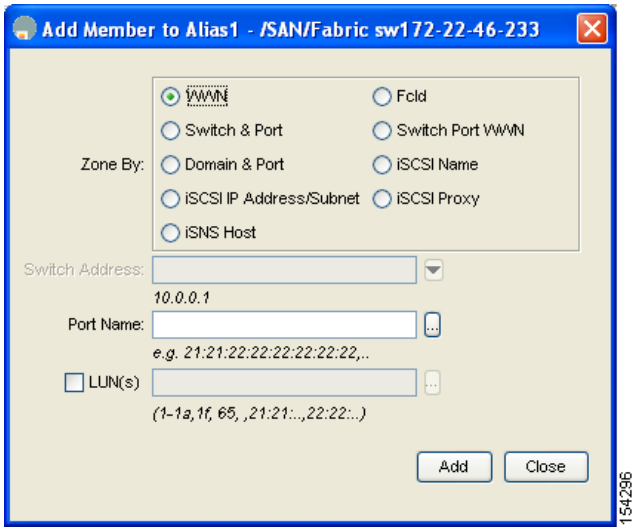

Step 4 Optionally, change **Zone By** to another criteria.

Step 5 Click**...** and select a port name to complete the Port Name field.

Step 6 Optionally, click**...** and select a LUN to complete the LUN(s) field.

Step 7 Click **Add** to add the member to the alias or click **Close** to close the dialog box.

Converting Zone members to pWWN-based Members

Fabric Manager Release 2.1(2) introduced the ability to convert zone and alias members from switch port or FC ID based membership to pWWN-based membership. You can use this feature to convert to pWWN so that your zone configuration does not change if a card or switch is changed in your fabric.

To convert switch port and FC ID members to pWWN members using Fabric Manager, follow these steps:

Step 1 Click Zone and then select Edit Local Full Zone Database.

You see the Select VSAN dialog box.

Step 2 Select the VSAN to be converted and click **OK**.

You see the zone information for that VSAN.

Step 3 Right-click any zone in the left pane and select **Convert Switch Port/FCID members to pWWN** (see Figure 26-19).

Figure 26-19 Convert Switch Port/FCID Members to pWWN

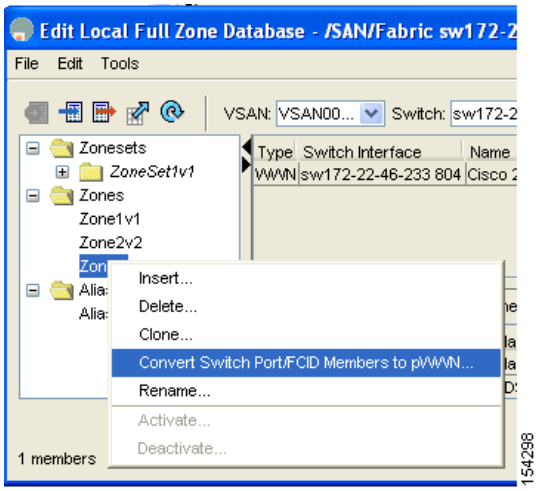

Step 4 You see the conversion dialog box, listing all members that will be converted.

Step 5 Verify the changes and click **Continue Conversion**.

Step 6 Click **Yes** in the confirmation dialog box to convert that member to pWWN-based membership.

Zone Sets

This section describes zone sets and includes the following topics:

About Zone Set Creation

In Figure 26-20, two separate sets are created, each with its own membership hierarchy and zone members.

Figure 26-20 Hierarchy of Zone Sets, Zones, and Zone Members

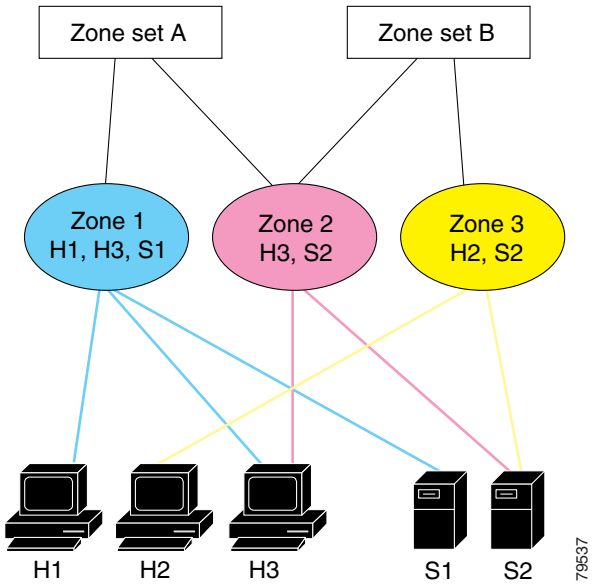

Zones provide a mechanism for specifying access control, while zone sets are a grouping of zones to enforce access control in the fabric. Either zone set A or zone set B can be activated (but not together).

Tip Zone sets are configured with the names of the member zones and the VSAN (if the zone set is in a configured VSAN).

Creating Zone Sets

To create a zone set to include several zones using Fabric Manager, follow these steps:

Step 1 Click **Zone** and then select **Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the drop-down menu.

If you chose **Zone > Edit Local Full Zone Database**, or the **All VSANs** folder is selected, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window shown in Figure 26-21 for the VSAN you selected.

Figure 26-21 Edit Local Full Zone Database

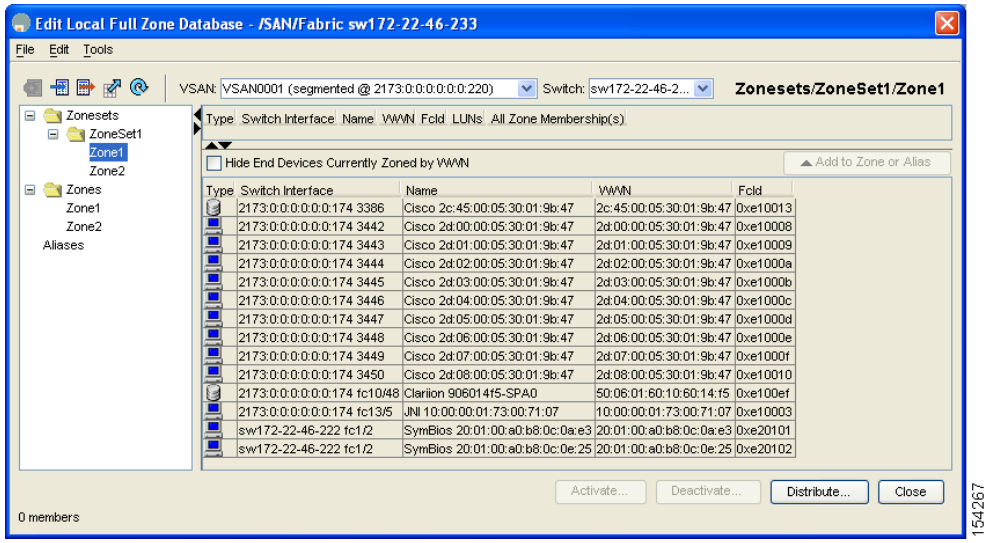

Step 2 Right-click the **Zonesets** folder in the left pane and select **Insert** to add a zone set (see Figure 26-22).

Figure 26-22 Right-click Zonesets Folder

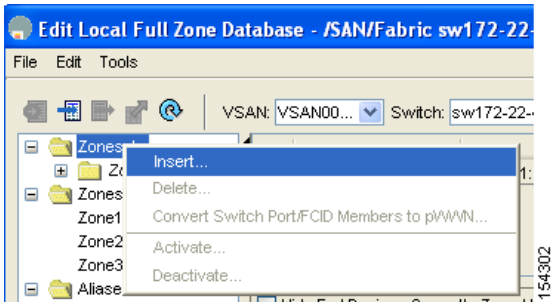

You see the Create Zoneset dialog box shown in Figure 26-23.

Figure 26-23 Create Zoneset

Step 3 Optionally, change the default name.

Step 4 Click OK.

The new zone set appears in the list in the left pane.

Step 5 Activate the zone set after creation by selecting it and then clicking **Activate** (see Figure 26-21). The configuration is distributed to the other switches in the network fabric.

Note When you confirm activation, the current running configuration is saved to the startup configuration. This permanently saves any changes made to the running configuration (not just zoning changes).

Adding Zones to a Zone Set

To add a zone to a zone set from the Edit Local Full Zone Database window, drag and drop the zone into the folder for the zone set.

Active and Full Zone Set Considerations

Before configuring a zone set, consider the following guidelines:

- Each VSAN can have multiple zone sets but only one zone set can be active at any given time.
- When you create a zone set, that zone set becomes a part of the full zone set.
- When you activate a zone set, a copy of the zone set from the full zone set is used to enforce zoning, and is called the active zone set.An active zone set cannot be modified. A zone that is part of an active zone set is called an active zone.
- The administrator can modify the full zone set even if a zone set with the same name is active. However, the modification will be enforced only upon reactivation.
- When the activation is done, the active zone set is automatically stored in persistent configuration. This enables the switch to preserve the active zone set information across switch resets.
- All other switches in the fabric receive the active zone set so they can enforce zoning in their respective switches.

- Hard and soft zoning are implemented using the active zone set. Modifications take effect during zone set activation.
- An FC ID or Nx port that is not part of the active zone set belongs to the default zone and the default zone information is not distributed to other switches.

Note If one zone set is active and you activate another zone set, the currently active zone set is automatically deactivated. You do not need to explicitly deactivate the currently active zone set before activating a new zone set.

Figure 26-24 shows active and full zone sets.

Figure 26-24 Active and Full Zone Sets

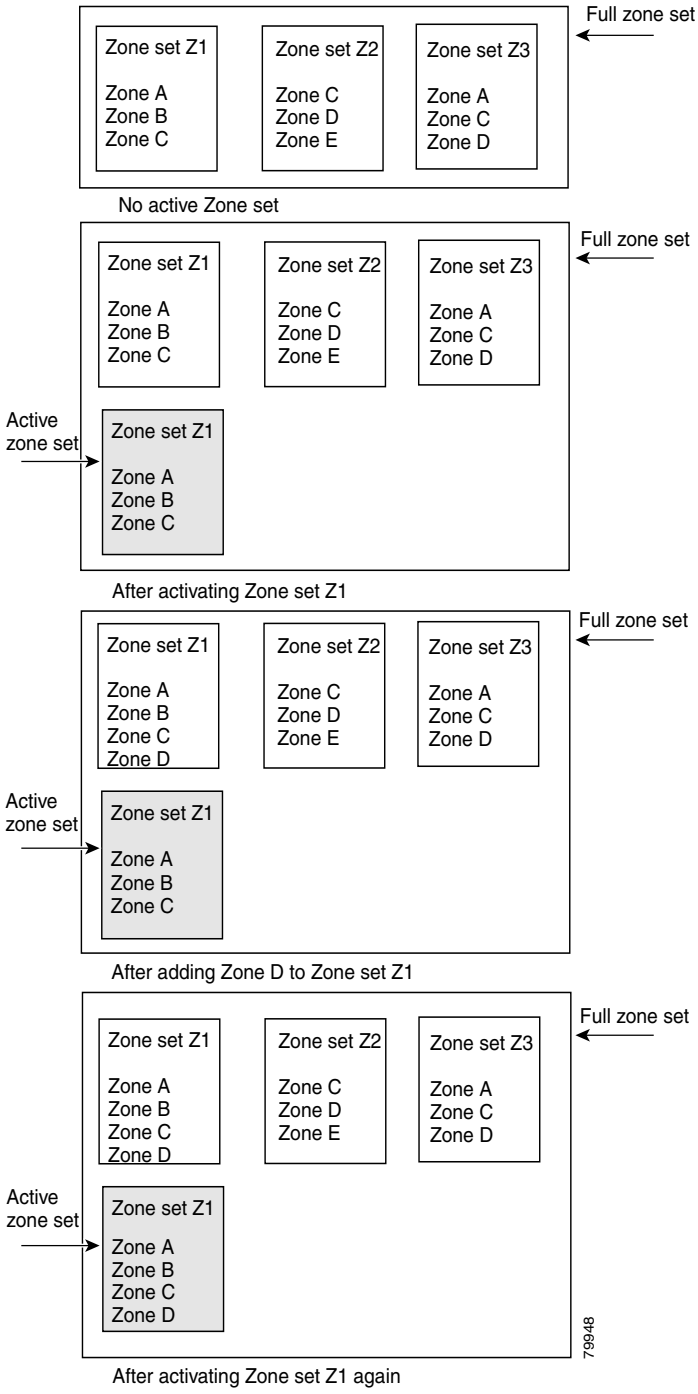

Activating or Deactivating a Zone Set

Changes to a zone set do not take effect in a full zone set until you activate it.

To activate or deactivate an existing zone set using Fabric Manager, follow these steps:

Step 1 Click **Zone** and then select **Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the drop-down menu.

If you chose **Zone > Edit Local Full Zone Database**, or the **All VSANs** folder is selected, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window for the VSAN you selected.

Step 2 Right-click the zone set folder in the left pane and then click **Activate** to activate the zone set or click **Deactivate** to deactivate the zone set.

You see the Save Configuration dialog box in Figure 26-25.

Figure 26-25 Save Configuration Options for a New Zone Set

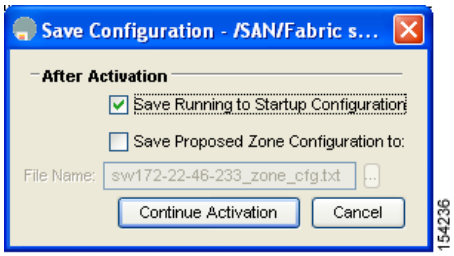

Step 3 If you are activating the zone set, then check the **Save Running to Startup Configuration** check box to save these changes to the startup configuration.

Step 4 Click **Continue Activation** to activate the zone set or **Yes** if you are deactivating the zone set (see Figure 26-25).

Note If one zone set is active and you activate another zone set, the currently active zone set is automatically deactivated.

Tip You do not have to copy the running configuration to the startup configuration to store the active zone set. However, you do need to copy the running configuration to the startup configuration to explicitly store full zone sets. It is not available across switch resets.

Caution If you deactivate the active zone set in a VSAN that is also configured for IVR, the active IVR zone set (IVZS) is also deactivated and all IVR traffic to and from the switch is stopped. This deactivation can disrupt traffic in more than one VSAN. Before deactivating the active zone set, check the active zone analysis for the VSAN. To reactivate the IVZS, you must reactivate the regular zone set (see the "Configuring IVR Zones and Zone Sets" section on page 25-23).

Caution If the currently active zone set contains IVR zones, activating the zone set from a switch where IVR is not enabled disrupts IVR traffic to and from that VSAN. We strongly recommend that you always activate the zone set from an IVR-enabled switch to avoid disrupting IVR traffic.

Zone Enforcement

Zoning can be enforced in two ways: soft and hard. Each end device (N port or NL port) discovers other devices in the fabric by querying the name server. When a device logs in to the name server, the name server returns the list of other devices that can be accessed by the querying device. If an Nx port does not know about the FC IDs of other devices outside its zone, it cannot access those devices.

In soft zoning, zoning restrictions are applied only during interaction between the name server and the end device. If an end device somehow knows the FC ID of a device outside its zone, it can access that device.

Hard zoning is enforced by the hardware on each frame sent by an Nx port. As frames enter the switch, source-destination IDs are compared with permitted combinations to allow the frame at wirespeed. Hard zoning is applied to all forms of zoning.

Note Hard zoning enforces zoning restrictions on every frame, and prevents unauthorized access.

Switches in the Cisco MDS 9000 Family support both hard and soft zoning.

About the Default Zone

Each member of a fabric (in effect a device attached to an Nx port) can belong to any zone. If a member is not part of any active zone, it is considered to be part of the default zone. Therefore, if no zone set is active in the fabric, all devices are considered to be in the default zone. Even though a member can belong to multiple zones, a member that is part of the default zone cannot be part of any other zone. The switch determines whether a port is a member of the default zone when the attached port comes up.

Note Unlike configured zones, default zone information is not distributed to the other switches in the fabric.

Traffic can either be permitted or denied among members of the default zone. This information is not distributed to all switches; it must be configured in each switch.

Note When the switch is initialized for the first time, no zones are configured and all members are considered to be part of the default zone. Members are not permitted to talk to each other.

Configure the default zone policy on each switch in the fabric. If you change the default zone policy on one switch in a fabric, be sure to change it on all the other switches in the fabric.

Note The default settings for default zone configurations can be changed.

The default zone members are explicitly listed when the default policy is configured as permit or when a zone set is active. When the default policy is configured as deny, the members of this zone are not explicitly enumerated when you view the active zone set.

You can change the default zone policy for any VSAN by choosing **VSAN***xx* > **Default Zone** from the Fabric Manager menu tree and clicking the **Policies** tab. It is recommended that you establish connectivity among devices by assigning them to a non-default zone.

The active zone set is shown in italic type. After you have made changes to the active zone set and before you activate the changes, the zone set is shown in boldface italic type.

Configuring the Default Zone

To permit or deny traffic to members in the default zone using Fabric Manager, follow these steps:

Step 1 Expand a **VSAN** and then select **Default Zone** in the Fabric Manager Logical Domains pane,

Step 2 Click the **Policies** tab in the Information pane.

You see the zone policies information in the Information pane (see Figure 26-26).

Figure 26-26 Default Zone Policies

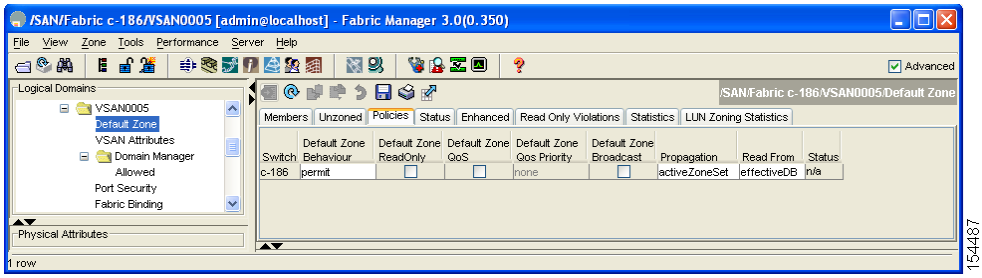

Step 3 Click the Default Zone Behavior field and choose either **permit** or **deny** from the pull-down menu (see Figure 26-26).

Zone Set Distribution

You can distribute full zone sets using one of two methods: one-time distribution or full zone set distribution. Table 26-1 lists the differences.

Table 26-1 Zone Set Distribution Differences

One-Time Distribution	Full Zone Set Distribution
Distributes the full zone set immediately.	Does not distribute the full zone set immediately.
Does not distribute the full zone set information along with the active zone set during activation, deactivation, or merge process.	Remembers to distribute the full zone set information along with the active zone set during activation, deactivation, and merge processes.

This section describes zone set distribution and includes the following topics:

Enabling Full Zone Set Distribution

All switches in the Cisco MDS 9000 Family distribute active zone sets when new E port links come up or when a new zone set is activated in a VSAN. The zone set distribution takes effect while sending merge requests to the adjacent switch or while activating a zone set.

To enable full zone set and active zone set distribution to all switches on a per VSAN basis using Fabric Manager, follow these steps:

Step 1 Expand a **VSAN** and select a zone set in the Logical Domains pane.

You see the zone set configuration in the Information pane. The Active Zones tab is the default.

Step 2 Click the **Policies** tab.

You see the configured policies for the zone in Figure 26-27.

Figure 26-27 Configured Policies for the Zone

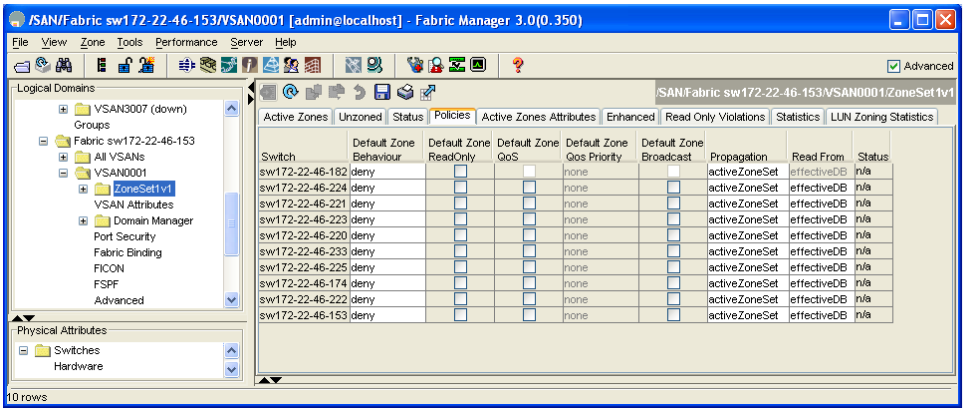

Step 3 Set the Propagation column to **fullZoneset** from the drop-down menu.

Step 4 Click **Apply Changes** to propagate the full zone set, or click **Undo Changes** to discard any changes you made.

Enabling a One-Time Distribution

You can perform a one-time distribution of inactive, unmodified zone sets throughout the fabric.

To propagate a one-time distribution of the full zone set using Fabric Manager, follow these steps:

Step 1 Click the **Zone** menu and then select **Edit Local Full Zone Database.**

Step 2 Select the appropriate zone from the list in the left pane.

You see the Edit Local Full Zone Database dialog box.

Step 3 Click **Distribute** to distribute the full zone set across the fabric.

This procedure only distributes the full zone set information—it does not save the information to the startup configuration. You must explicitly save the running configuration to the startup configuration to save the full zone set information to the startup configuration.

Note The one-time distribution of the full zone set is supported in interop 2 and interop 3 modes—not in interop 1 mode.

About Recovering from Link Isolation

When two switches in a fabric are merged using a TE or E port, these TE and E ports may become isolated when the active zone set databases are different between the two switches or fabrics. When a TE port or an E port become isolated, you can recover that port from its isolated state using one of three options:

- Import the neighboring switch's active zone set database and replace the current active zone set (see Figure 26-28).
- Export the current database to the neighboring switch (see Figure 26-28).
- Manually resolve the conflict by editing the full zone set, activating the corrected zone set, and then bringing up the link.

Figure 26-28 Importing and Exporting the Database

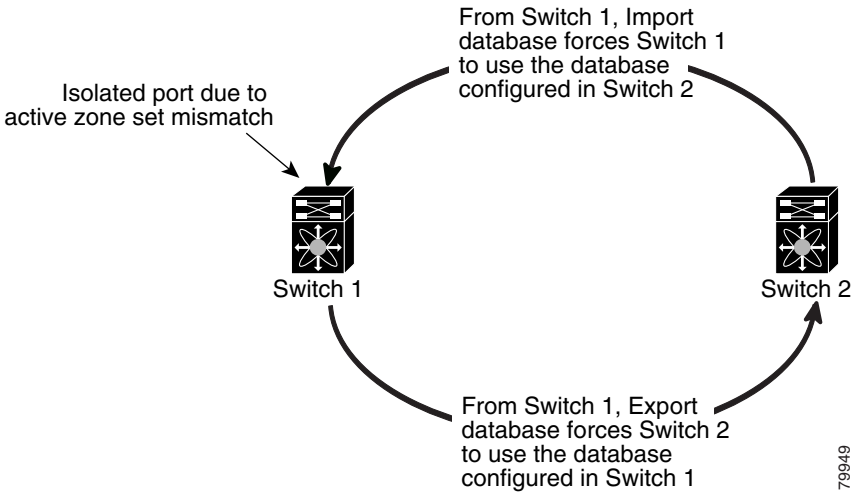

Importing and Exporting Zone Sets

To import or export the zone set information from or to an adjacent switch using Fabric Manager, follow these steps:

Step 1 Click **Zone** in the Zone menu and then select **Merge Fail Recovery**.

You see the Zone Merge Failure Recovery dialog box in Figure 26-29.

Figure 26-29 Zone Merge Fail Recovery

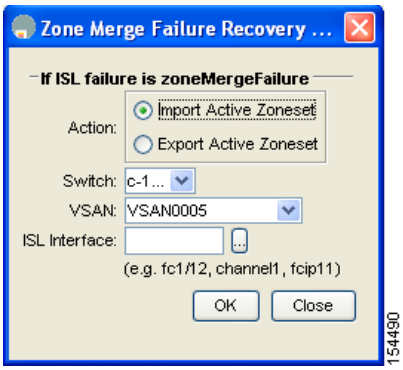

Step 2 Select the **Import Active Zoneset** radio button or the **Export Active Zoneset** radio button.

Step 3 Select the switch from which to import or export the zone set information from the drop-down list.

Step 4 Select the VSAN from which to import or export the zone set information from the drop-down list.

Step 5 Select the interface to use for the import process.

Step 6 Click **OK** to import or export the active zone set, or click **Close** to close the dialog box without importing or exporting the active zone set.

Note Issue the import and export from a single switch. Importing from one switch and exporting from another switch can lead to isolation again.

Zone Set Duplication

You can make a copy and then edit it without altering the existing active zone set. You can copy an active zone set from the bootflash: directory, volatile: directory, or slot0, to one of the following areas:

- To the full zone set
- To a remote location (using FTP, SCP, SFTP, or TFTP).

The active zone set is not part of the full zone set. You cannot make changes to an existing zone set and activate it, if the full zone set is lost or is not propagated.

Caution Copying an active zone set to a full zone set may overwrite a zone with the same name, if it already exists in the full zone set database.

This section includes the following topics:

Copying Zone Sets

On the Cisco MDS Family switches, you cannot edit an active zone set. However, you can copy an active zone set to create a new zone set that you can edit.

To make a copy of a zone set using Fabric Manager, follow these steps:

Step 1 Click **Zone** and then select **Copy Full Zone Database** from the Zone menu.

You see the Copy Full Zone Database dialog box in Figure 26-30.

Figure 26-30 Copy Full Zone Database

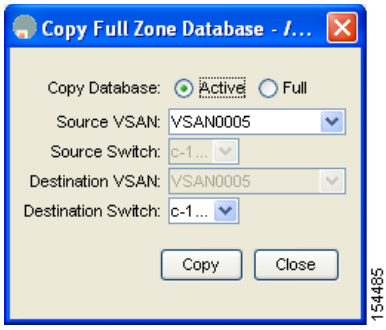

Step 2 Select the **Active** or the **Full** radio button, depending on which type of database you want to copy.

Step 3 Select the source VSAN from the drop-down list.

Step 4 If you selected **Copy Full**, select the source switch and the destination VSAN from those drop-down lists.

Step 5 Select the destination switch from the drop-down list.

Step 6 Click **Copy** to copy the database, or click **Close** to close the dialog box without copying.

Caution If the Inter-VSAN Routing (IVR) feature is enabled and if IVR zones exist in the active zone set, then a zone set copy operation copies all the IVR zones to the full zone database. To prevent copying to the IVR zones, you must explicitly remove them from the full zone set database before performing the copy operation. Refer to the Chapter 25, "Configuring Inter-VSAN Routing" for more information on the IVR feature.

Renaming Zones, Zone Sets, and Aliases

To rename a zone, zone set, or alias using Fabric Manager, follow these steps:

Step 1 Choose **Zone > Edit Local Full Zone Database** from the Zone menu shown in Figure 26-31 or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.

If you chose **Zone > Edit Local Full Zone Database**, or the **All VSANs** folder is selected, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window for the VSAN you selected (see Figure 26-31).

Figure 26-31 Edit Local Full Zone Database

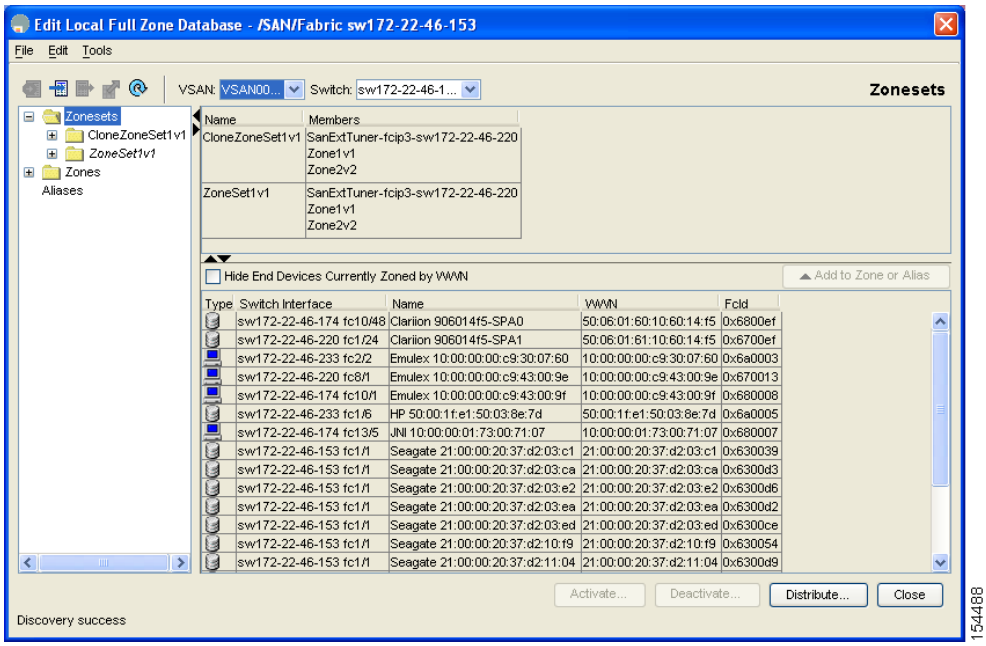

Step 2 Select a zone.

Step 3 Click **Edit > Rename** from the menu.

An edit box appears around the zone name.

Step 4 Supply a new name.

Step 5 Click **Activate** or **Distribute**.

Cloning Zones, Zone Sets, fcaliases, and Zone Attribute Groups

To clone a zone, zone set, fcalias, or zone-attribute-group, follow these steps:

Step 1 Choose **Zone > Edit Local Full Zone Database** from the Zone menu or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.

If you chose **Zone > Edit Local Full Zone Database**, or the **All VSANs** folder is selected, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window for the VSAN you selected.

Step 2 Click **Edit > Clone** from the dialog box menu.

You see the Clone Zoneset dialog box. The default name is the word **Clone** followed by the original name.

Figure 26-32 Clone Zoneset Dialog Box

Step 3 Optionally, change the name for the cloned entry.

Step 4 Click **OK** to save the new clone or click **Cancel** to close the dialog box and discard any unsaved changes.

The cloned database now appears along with the original database.

About Backing Up and Restoring Zones

You can back up the zone configuration to a workstation using TFTP. This zone backup file can then be used to restore the zone configuration on a switch. Restoring the zone configuration overwrites any existing zone configuration on a switch.

Backing Up and Restoring Zones

To back up or restore the full zone configuration using Fabric Manager, follow these steps:

Step 1 Click **Zone > Edit Local Full Zone Database** from the Zone menu or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.

If you chose **Zone > Edit Local Full Zone Database**, or the **All VSANs** folder is selected, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window for the VSAN you selected.

Step 2 Choose **File > Backup** to back up the existing zone configuration to a workstation using TFTP.

Step 3 Click **File > Restore** to restore a saved zone configuration. You can optionally edit this configuration before restoring it to the switch.

Migrating a Non-MDS Database

To use the Zone Migration Wizard to migrate a non-MDS database using Fabric Manager, follow these steps:

Step 1 Click **Zone** and select **Migrate Non-MDS Database** from the Zone menu.

You see the Zone Migration Wizard.

Step 2 Follow the prompts in the wizard to migrate the database.

Clearing the Zone Server Database

You can clear all configured information in the zone server database for the specified VSAN. To clear the zone server database, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

Note Clearing a zone set only erases the full zone database, not the active zone database.

Note After clearing the zone server database, you must explicitly copy the running configuration to the startup configuration to ensure that the running configuration is used when the switch reboots.

Advanced Zone Attributes

This section describes advanced zone attributes and includes the following topics:

About Zone-Based Traffic Priority

The zoning feature provides an additional segregation mechanism to prioritize select zones in a fabric and set up access control between devices. Using this feature, you can configure the Quality of Service (QoS) priority as a zone attribute. You can assign the QoS traffic priority attribute to be high, medium, or low. By default, zones with no specified priority are implicitly assigned a low priority. See the "VSAN Versus Zone-Based QoS" section on page 60-6 for more information.

To use this feature, you need to obtain the ENTERPRISE_PKG license (see Chapter 10, "Obtaining and Installing Licenses") and you must enable QoS in the switch (see the "About Class Map Creation" section on page 60-7).

This feature allows SAN administrators to configure QoS in terms of a familiar data flow identification paradigm. You can configure this attribute on a zone-wide basis rather than between zone members.

Caution If zone-based QoS is implemented in a switch, you cannot configure the interop mode in that VSAN.

Configuring Zone-Based Traffic Priority

To configure zone priority using Fabric Manager, follow these steps:

Step 1 Expand a **VSAN** and then select a zone set in the Logical Domains pane.

Step 2 Click the **Policies** tab in the Information pane.

You see the Zone policy information in the Information pane in Figure 26-33.

Figure 26-33 Zone Policies Tab in the Information Pane

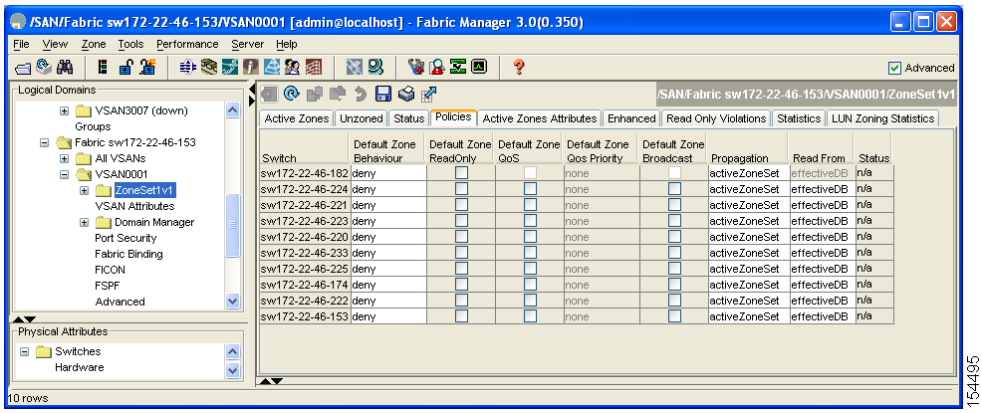

Step 3 Use the check boxes and drop-down menus to configure QoS on the default zone (see Figure 26-33).

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard these changes.

Configuring Default Zone QoS Priority Attributes

QoS priority attribute configuration changes take effect when you activate the zone set of the associated zone.

Note If a member is part of two zones with two different QoS priority attributes, the higher QoS value is implemented. This situation does not arise in the VSAN-based QoS as the first matching entry is implemented.

To configure the QoS priority attributes for a default zone using Fabric Manager, follow these steps:

Step 1 Choose **Zone > Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.

If you chose **Zone > Edit Local Full Zone Database**, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window for the VSAN you selected.

Step 2 Click the **Edit** menu and select **Edit Default Zone Attributes** to configure the default zone QoS priority attributes.

Figure 26-34 Edit Local Full Zone Database Window

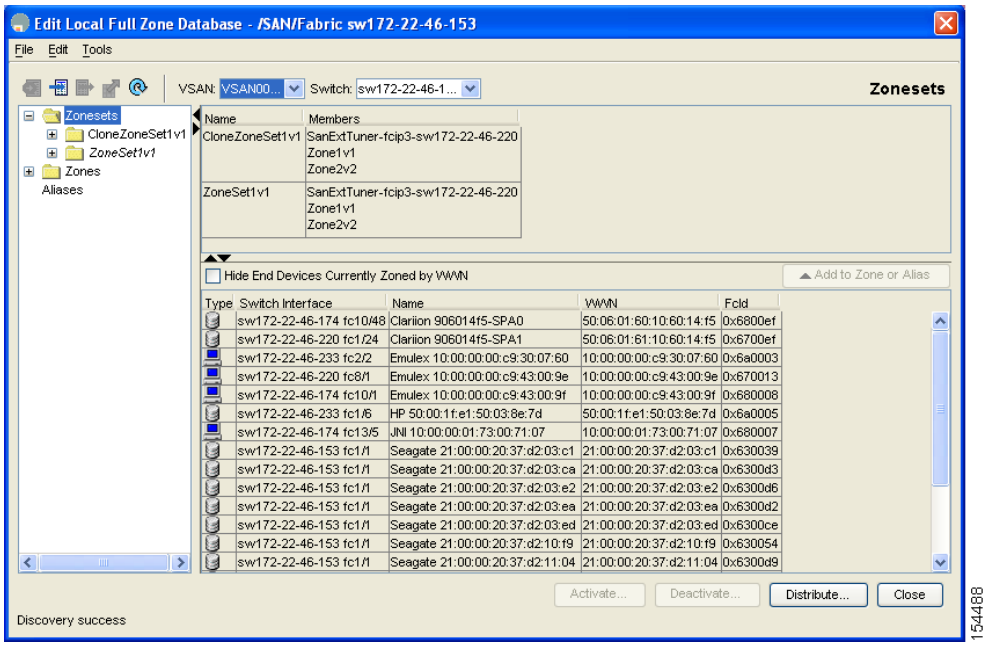

Step 3 Check the **Permit QoS Traffic with Priority** check box and set the Qos priority drop-down menu to **low**, **medium**, or **high**.

Step 4 Click **OK** to save these changes or click **Close** to discard any unsaved changes.

Configuring the Default Zone Policy

To permit or deny traffic in the default zone using Fabric Manager, follow these steps:

Step 1 Choose **Zone > Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.

If you chose **Zone > Edit Local Full Zone Database**, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window for the VSAN you selected.

Figure 26-35 Edit Local Full-Zone Database

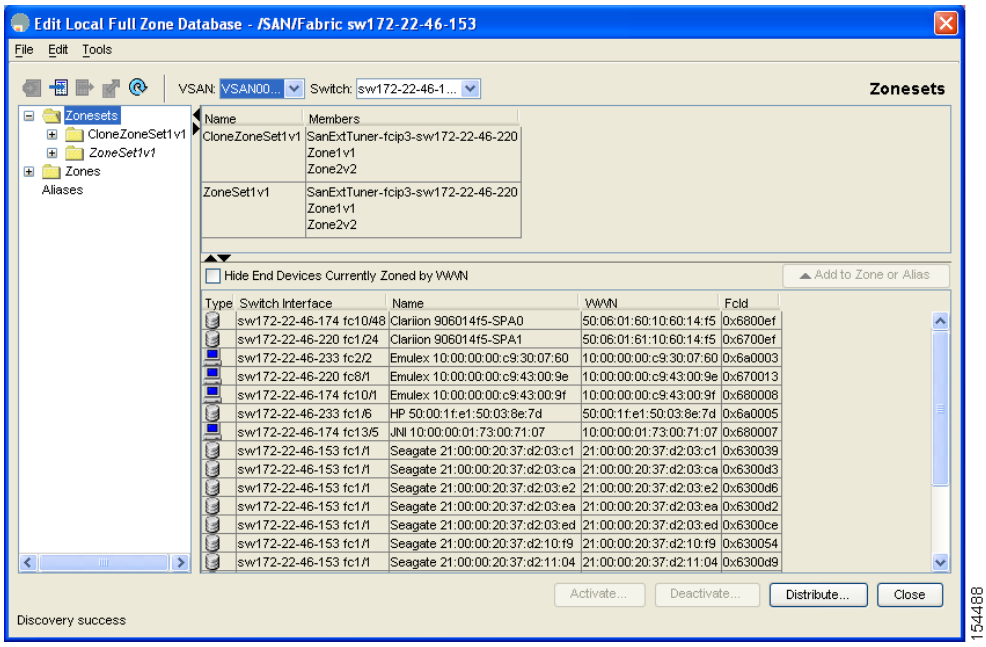

Step 2 Click the **Edit** menu and select **Edit Default Zone Attributes** to configure the default zone QoS priority attributes.

You see the Modify Default Zone Properties dialog box shown in Figure 26-36.

Figure 26-36 Modify Default Zone Properties Dialog Box

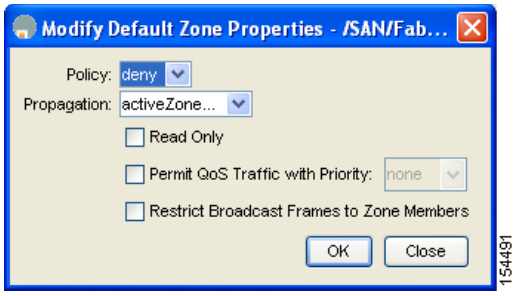

Step 3 Set the Policy drop-down menu to **permit** to permit traffic in the default zone, or set it to **deny** to block traffic in the default zone (see Figure 26-36).

Step 4 Click **OK** to save these changes or click **Close** to discard any unsaved changes (see Figure 26-36).

About Broadcast Zoning

You can configure broadcast frames in the basic zoning mode. By default, broadcast zoning is disabled and broadcast frames are sent to all Nx ports in the VSAN. When enabled, broadcast frames are only sent to Nx ports in the same zone, or zones, as the sender. Enable broadcast zoning when when a host or storage device uses this feature.

Table 26-2 identifies the rules for the delivery of broadcast frames.

Table 26-2 Broadcasting Requirements

Active Zoning?	Broadcast Enabled?	Frames Broadcast?	Comments
Yes	Yes	Yes	Broadcast to all Nx ports that share a broadcast zone with the source of broadcast frames.
No	Yes	Yes	Broadcast to all Nx ports.
Yes	No	No	Broadcasting is disabled.

Tip If any NL port attached to an FL port shares a broadcast zone with the source of the broadcast frame, then the frames are broadcast to all devices in the loop.

Caution If broadcast zoning is enabled on a switch, you cannot configure the interop mode in that VSAN.

Configuring Broadcast Zoning

To broadcast frames in the basic zoning mode using Fabric Manager, follow these steps:

Step 1 Expand a **VSAN** and then select a zone set in the Logical Domains pane.

Step 2 Click the **Policies** tab in the Information pane.

You see the Zone policy information in the Information pane in Figure 26-37.

Figure 26-37 Zone Policy Information

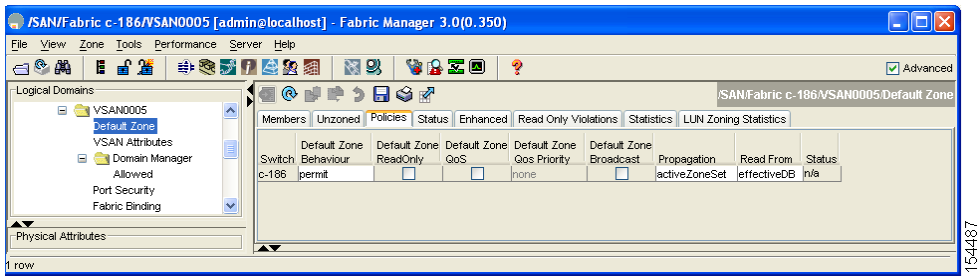

Step 3 Check the **Broadcast** check box to enable broadcast frames on the default zone (see Figure 26-37).

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard these changes (see Figure 26-37).

About LUN Zoning

Logical unit number (LUN) zoning is a feature specific to switches in the Cisco MDS 9000 Family.

Caution LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure the interop mode in that switch.

A storage device can have multiple LUNs behind it. If the device port is part of a zone, a member of the zone can access any LUN in the device. With LUN zoning, you can restrict access to specific LUNs associated with a device.

Note When LUN 0 is not included within a zone, then, as per standards requirements, control traffic to LUN 0 (for example, REPORT_LUNS, INQUIRY) is supported, but data traffic to LUN 0 (for example, READ, WRITE) is denied.

- Host H1 can access LUN 2 in S1 and LUN 0 in S2. It cannot access any other LUNs in S1 or S2.
- Host H2 can access LUNs 1 and 3 in S1 and only LUN 1 in S2. It cannot access any other LUNs in S1 or S2.

Note Unzoned LUNs automatically become members of the default zone.

Figure 26-38 shows a LUN-based zone example.

Figure 26-38 LUN Zoning Access

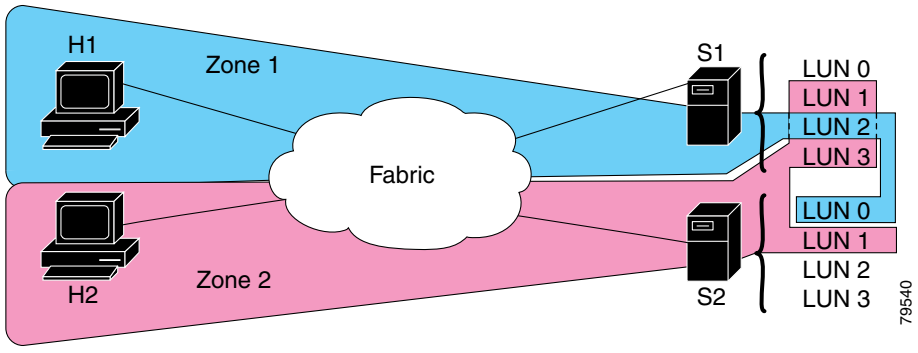

Configuring a LUN-Based Zone

To configure a LUN-based zone using Fabric Manager, follow these steps:

Step 1 Choose **Zone > Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.

If you chose **Zone > Edit Local Full Zone Database**, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window for the VSAN you selected (see Figure 26-39).

Figure 26-39 Edit Local Full Zone Database Window

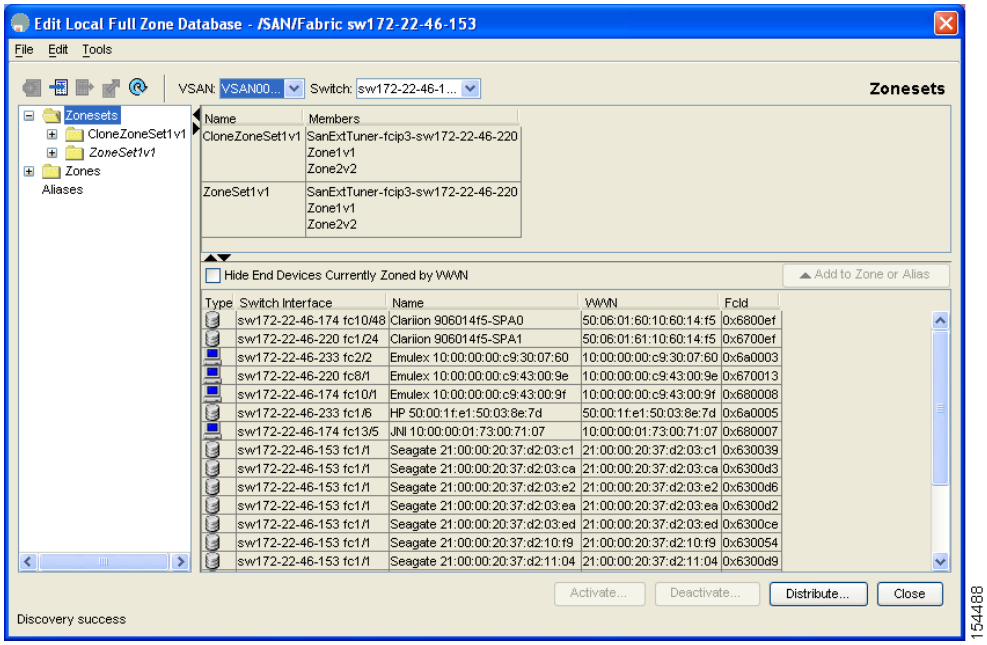

Step 2 Right-click the **Zones** folder in the left pane and then select **Insert** to add a zone (see Figure 26-40).

Figure 26-40 Right-Click the Zones Folder and Select Insert to Create a New Zone

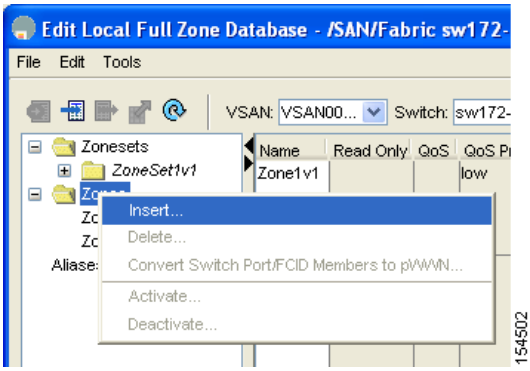

You see the Insert Zone dialog box shown in Figure 26-41.

Figure 26-41 Insert Zone Dialog Box

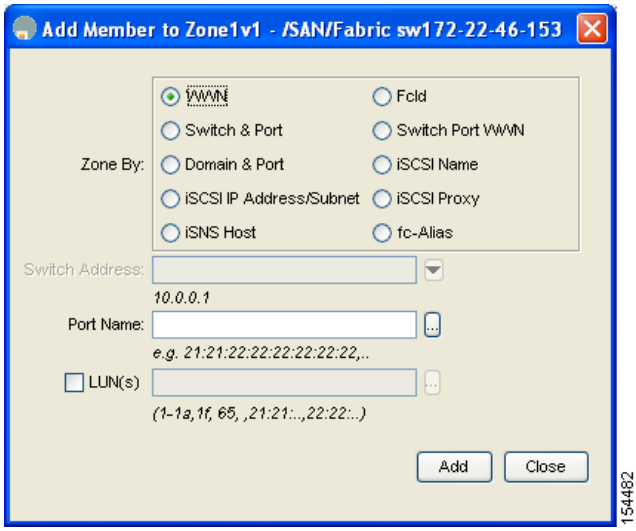

Step 3 Select either **WWN** or **FCID** radio button for the Zone By options to create a LUN-based zone.

Step 4 Check the **LUN** check box and add the LUNs for this zone in the text box.

Step 5 Click **Add** to add this LUN-based zone or **Close** to close the dialog box without adding the LUN-based zone.

Assigning LUNs to Storage Subsystems

LUN masking and mapping restricts server access to specific LUNs. If LUN masking is enabled on a storage subsystem and if you want to perform additional LUN zoning in a Cisco MDS 9000 Family switch, obtain the LUN number for each host bus adapter (HBA) from the storage subsystem and then configure the LUN-based zone procedure provided in the "Configuring a LUN-Based Zone" section on page 26-39.

Note Refer to the relevant user manuals to obtain the LUN number for each HBA.

Caution If you make any errors when configuring this scenario, you are prone to lose data.

About Read-Only Zones

By default, an initiator has both read and write access to the target's media when they are members of the same Fibre Channel zone. The read-only zone feature allows members to have only read access to the media within a read-only Fibre Channel zone.

You can also configure LUN zones as read-only zones.

Any zone can be identified as a read-only zone. By default all zones have read-write permission unless explicitly configured as a read-only zone.

Follow these guidelines when configuring read-only zones:

- If read-only zones are implemented, the switch prevents write access to user data within the zone.
- If two members belong to a read-only zone and to a read-write zone, read-only zone has priority and write access is denied.
- LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure interop mode in that switch.
- Read-only volumes are not supported by some operating system and file system combinations (for example, Windows NT or Windows 2000 and NTFS file system). Volumes within read-only zones are not available to such hosts. However, if these hosts are already booted when the read-only zones are activated, then read-only volumes are available to those hosts.

  The read-only zone feature behaves as designed if FAT16 or FAT32 file system is used with the previously-mentioned Windows operating systems.

Configuring Read-Only Zones

To configure read-only zones using Fabric Manager, follow these steps:

Step 1 Choose **Zone > Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.

If you chose **Zone > Edit Local Full Zone Database**, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window for the VSAN you selected.

Step 2 Right-click the **Zones** folder in the Edit Local Full Zone Database dialog box for that VSAN and select **Insert** to add a zone.

Step 3 You see the Create Zone Dialog Box.

Figure 26-42 Create Zone Dialog Box

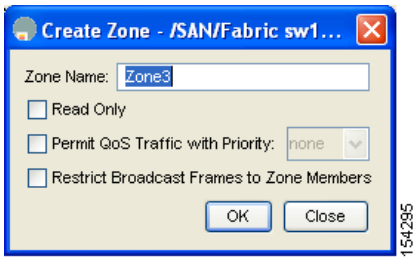

Step 4 Check the **Read Only** check box to create a read-only zone.

Step 5 Click **OK**.

To configure the read-only option for a default zone, follow these steps

Step 1 Choose **Zone > Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.

If you chose **Zone > Edit Local Full Zone Database**, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window for the VSAN you selected.

Step 2 Click **Edit > Edit Default Zone Attributes** to configure the default zone read-only attributes as shown in Figure 26-43.

Figure 26-43 Edit Default Zone Attributes

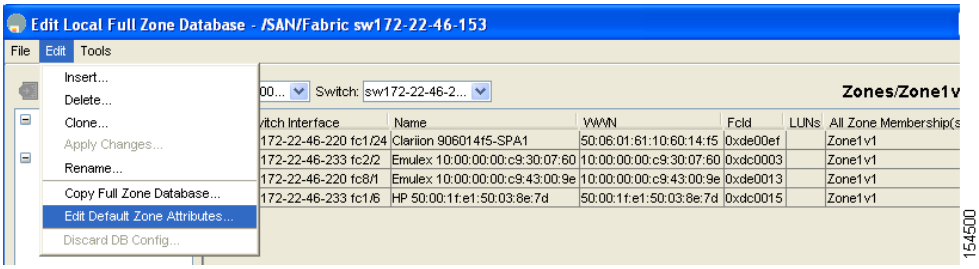

You see the Modify Default Zone Properties dialog box.

Step 3 Check the **Read Only** check box.

Step 4 Click **OK** to save these changes or click **Close** to discard any unsaved changes.

Displaying Zone Information

To view zone information and statistics using Fabric Manager, follow these steps:

Step 1 Expand a **VSAN** and select a zone set in the Logical Domains pane.

You see the zone configuration in the Information pane.

Step 2 Choose the **Read Only Violations**, **Statistics** tab, or **LUN Zoning Statistics** tab to view statistics for the selected zone.

Enhanced Zoning

The zoning feature complies with the FC-GS-4 and FC-SW-3 standards. Both standards support the basic zoning functionalities explained in the previous section and the enhanced zoning functionalities described in this section.

This section includes the following topics:

About Enhanced Zoning

Table 26-3 lists the advantages of the enhanced zoning feature in all switches in the Cisco MDS 9000 Family.

Table 26-3 Advantages of Enhanced Zoning

Basic Zoning	Enhanced Zoning	Enhanced Zoning Advantages
Administrators can make simultaneous configuration changes. Upon activation, one administrator can overwrite another administrator's changes.	Performs all configurations within a single configuration session. When you begin a session, the switch locks the entire fabric to implement the change.	One configuration session for the entire fabric to ensure consistency within the fabric.
If a zone is part of multiple zone sets, you create an instance of this zone in each zone set	References to the zone are used by the zone sets as required once you define the zone.	Reduced payload size as the zone is referenced. The size is more pronounced with bigger databases.
The default zone policy is defined per switch. To ensure smooth fabric operation, all switches in the fabric must have the same default zone setting.	Enforces and exchanges the default zone setting throughout the fabric.	Fabric-wide policy enforcement reduces troubleshooting time.
To retrieve the results of the activation on a per switch basis, the managing switch provides a combined status about the activation. It does not identify the failure switch.	Retrieves the activation results and the nature of the problem from each remote switch.	Enhanced error reporting eases the troubleshooting process
To distribute the zoning database, you must reactivate the same zone set. The reactivation may affect hardware changes for hard zoning on the local switch and on remote switches. **Note** The size of the zoning database is 2 MB per VSAN.	Implements changes to the zoning database and distributes it without reactivation.	Distribution of zone sets without activation avoids hardware changes for hard zoning in the switches.

Table 26-3 Advantages of Enhanced Zoning (continued)

Basic Zoning	Enhanced Zoning	Enhanced Zoning Advantages
The MDS-specific zone member types (IP address, symbolic node name, and other types) may be used by other non-Cisco switches. During a merge, the MDS-specific types can be misunderstood by the non-Cisco switches.	Provides a vendor ID along with a vendor-specific type value to uniquely identify a member type.	Unique vendor type.
The fWWN-based zone membership is only supported in Cisco interop mode.	Supports fWWN-based membership in the default interop mode.	The fWWN-based member type is standardized.

Changing from Basic Zoning to Enhanced Zoning

To change to the enhanced zoning mode from the basic mode, follow these steps:

Step 1 Verify that all switches in the fabric are capable of working in the enhanced mode.

If one or more switches are not capable of working in enhanced mode, then your request to move to enhanced mode is rejected.

Step 2 Set the operation mode to enhanced zoning mode. By doing so, you will automatically start a session, acquire a fabric wide lock, distribute the active and full zoning database using the enhanced zoning data structures, distribute zoning policies and then release the lock. All switches in the fabric then move to the enhanced zoning mode.

Tip After moving from basic zoning to enhanced zoning we recommend that you save the running configuration.

Changing from Enhanced Zoning to Basic Zoning

The standards do not allow you to move back to basic zoning. However, Cisco MDS switches allow this move to enable you to downgrade and upgrade to other Cisco SAN-OS releases.

To change to the basic zoning mode from the enhanced mode, follow these steps:

Step 1 Verify that the active and full zone set do not contain any configuration that is specific to the enhanced zoning mode.

If such configurations exist, delete them before proceeding with this procedure. If you do not delete the existing configuration, the Cisco SAN-OS software automatically removes them.

Step 2 Set the operation mode to basic zoning mode. By doing so, you will automatically start a session, acquire a fabric wide lock, distribute the zoning information using the basic zoning data structure, apply the configuration changes and release the lock from all switches in the fabric. All switches in the fabric then move to basic zoning mode.

Note If a switch running Cisco MDS SAN-OS Release 2.0(1b), or later, with enhanced zoning enabled is downgraded to Cisco MDS SAN-OS Release 1.3(4), or earlier, the switch comes up in basic zoning mode and thus cannot join the fabric because all the other switches in the fabric are still in enhanced zoning mode.

Enabling Enhanced Zoning

By default, the enhanced zoning feature is disabled in all switches in the Cisco MDS 9000 Family.

To enable enhanced zoning in a VSAN using Fabric Manager, follow these steps

Step 1 Expand a VSAN and then select a zone set in the Logical Domains pane.

You see the zone set configuration in the Information pane.

Step 2 Click the **Enhanced** tab.

You see the current enhanced zoning configuration.

Step 3 Set the Action drop-down menu to **enhanced** to enable enhanced zoning in this VSAN.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard these changes.

Creating Attribute Groups

In enhanced mode, you can directly configure attributes using attribute groups.

To configure attribute groups, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

About Merging the Database

The merge behavior depends on the fabric-wide merge control setting:

- Restrict—If the two database are not identical, the ISLs between the switches are isolated.
- Allow—The two databases are merged using the merge rules specified in Table 26-4.

Table 26-4 Database Zone Merge Status

Local Database	Adjacent Database	Merge Status	Results of the Merge
The databases contain zone sets with the same name[1] but different zones, aliases, and attributes groups.		Successful.	The union of the local and adjacent databases.
The databases contains a zone, zone alias, or zone attribute group object with same name[1] but different members.		Failed.	ISLs are isolated.
Empty.	Contains data.	Successful.	The adjacent database information populates the local database.
Contains data.	Empty.	Successful.	The local database information populates the adjacent database.

1. In the enhanced zoning mode, the active zone set does not have a name in interop mode 1. The zone set names are only present for full zone sets.

The Merge Process includes the following steps:

1. The software compares the protocol versions. If the protocol versions differ, then the ISL is isolated.
2. If the protocol versions are the same, then the zone policies are compared. If the zone policies differ, then the ISL is isolated.
3. If the zone merge options are the same, then the comparison is implemented based on the merge control setting.
   a. If the setting is restrict, the active zone set and the full zone set should be identical. Otherwise the link is isolated.
   b. If the setting is allow, then the merge rules are used to perform the merge (see Table 26-4).

Analyzing a Zone Merge

To perform a zone merge analysis using Fabric Manager, follow these steps:

Step 1 Choose **Zone > Merge Analysis** from the Zone menu.

You see the Zone Merge Analysis dialog box shown in Figure 26-44.

Figure 26-44 Zone Merge Analysis Dialog Box

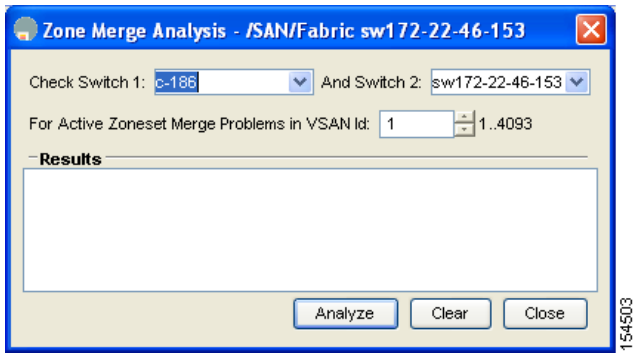

Step 2 Select the first switch to be analyzed from the Check Switch 1 drop-down list.

Step 3 Select the second switch to be analyzed from the And Switch 2 drop-down list.

Step 4 Enter the VSAN ID where the zone set merge failure occurred in the For Active Zoneset Merge Problems in VSAN Id field.

Step 5 Click **Analyze** to analyze the zone merge. Click **Clear** to clear the analysis data from the Zone Merge Analysis dialog box.

Configuring Zone Merge Control Policies

To configure merge control policies, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

Compacting the Zone Database for Downgrading

To compact the zone database for downgrading, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

Prior to Cisco MDS SAN-OS Release 3.0(1), only 2000 zones were supported per VSAN. If more than 2000 zones are added then a configuration check is registered to indicate that downgrading to a previous release could cause you to lose the zones over the limit. To avoid the configuration check you delete the excess zones and compact the zone database for the VSAN. If there are 2000 zones or fewer after the delete excess zones, the compacting process reissues an zone IDs and the configuration can be supported by previous versions.

Note To successfully downgrade, make sure that the zone database has 2000 or fewer zones.

Default Settings

Table 26-5 lists the default settings for basic zone parameters.

Table 26-5 Default Basic Zone Parameters

Parameters	Default
Default zone policy	Denied to all members.
Full zone set distribute	The full zone set(s) is not distributed.
Zone based traffic priority	Low.
Read-only zones	Read-write attributes for all zones.
Broadcast frames	Sent to all Nx ports.
Broadcast zoning	Disabled.
Enhanced zoning	Disabled.

CHAPTER 27

Distributing Device Alias Services

All switches in the Cisco MDS 9000 Family support Distributed Device Alias Services (device alias) on a per VSAN basis and on a fabric-wide basis. Device alias distribution allows you to move host bus adapters (HBAs) between VSANs without manually reentering alias names.

This chapter includes the following sections:

About Device Aliases

When the port WWN of a device must be specified to configure different features (zoning, QoS, port security) in a Cisco MDS 9000 Family switch, you must assign the right device name each time you configure these features. An inaccurate device name may cause unexpected results. You can circumvent this problem if you define a user-friendly name for a port WWN and use this name in all the configuration commands as required. These user-friendly names are referred to as *device aliases* in this chapter.

This section includes the following topics:

Device Alias Features

Device aliases have the following features:

- The device alias information is independent of your VSAN configuration.
- The device alias configuration and distribution is independent of the zone server and the zone server database.
- You can import legacy zone alias configurations without loosing data.

- The device alias application uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management and distribution. Device aliases use the coordinated distribution mode and the fabric-wide distribution scope (see Chapter 12, "Using the CFS Infrastructure").
- When you configure zones, IVR zones, or QoS features using device aliases, and if you display these configurations, you see that the device aliases are displayed along with their respective pWWNs.

Device Alias Requirements

Device aliases have the following requirements:

- You can only assign device aliases to pWWNs.
- The mapping between the pWWN and the device alias to which it is mapped must have a one-to-one relationship. A pWWN can be mapped to only one device alias and vice versa.
- A device alias name is restricted to 64 alphanumeric characters and may include one or more of the following characters:
  - a to z and A to Z
  - 1 to 9
  - - (hyphen) and _ (underscore)
  - $ (dollar sign) and ^ (up carat)

Zone Aliases Versus Device Aliases

Table 27-1 compares the configuration differences between zone-based alias configuration and device alias configuration.

Table 27-1 Comparison Between Zone Aliases and Device Aliases

Zone-Based Aliases	Device Aliases
Aliases are limited to the specified VSAN.	You can define device aliases without specifying the VSAN number. You can also use the same definition in one or more VSANs without any restrictions.
Zone aliases are part of the zoning configuration; the alias mapping cannot be used to configure other features.	Device aliases can be used with any feature that uses the pWWN.
You can use any zone member type to specify the end devices.	Only pWWNs are supported along with new device aliases like IP addresses.
Configuration is contained within the zone server database and is not available to other features.	Device aliases are not restricted to zoning. Device alias configuration is available to the FCNS, zone, fcping, traceroute, and IVR applications.

Device Alias Databases

The device alias feature uses two databases to accept and implement device alias configurations.

- Effective database—The database currently used by the fabric.
- Pending database—Your subsequent device alias configuration changes are stored in the pending database.

If you modify the device alias configuration, you need to commit or discard the changes as the fabric remains locked during this period.

This section includes the following topics:

About Device Alias Distribution

By default, device alias distribution is enabled. The device alias feature uses the coordinated distribution mechanism to distribute the modifications to all switches in a fabric.

If you have not committed the changes and you enable distribution, then a commit task will fail.

Distributing the Device Alias Database

To enable the device alias distribution using Fabric Manager, follow these steps:

Step 1 Expand **End Devices** and then select **Device Alias** in the Physical Attributes pane.

You see the device alias configuration in the Information pane Figure 27-1.

Figure 27-1 Device Aliases in Fabric Manager

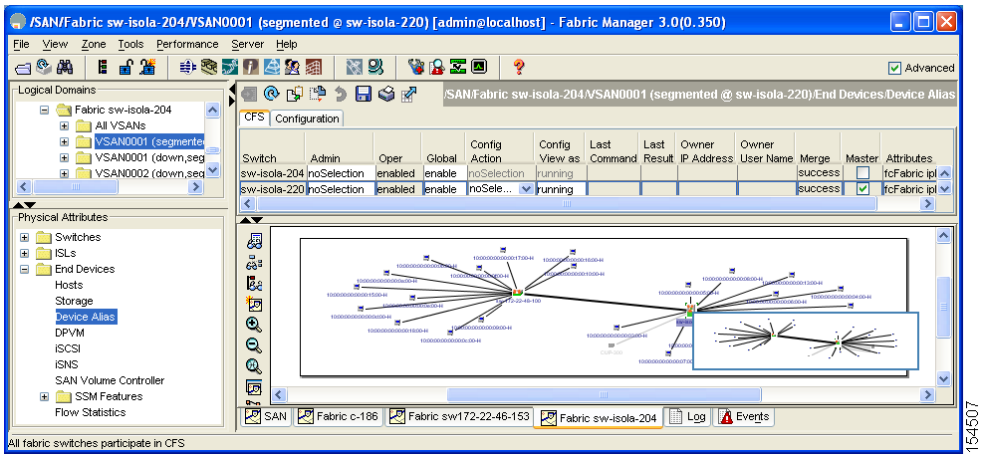

The **CFS** tab is the default tab.

Step 2 Select **enable** from the Global drop-down menus to enabled switch aliases.

Step 3 Select **commit** from the Config Action drop-down menu for the newly enabled switches.

Step 4 Click **Apply Changes** to commit and distribute these changes or click **Undo Changes** to discard any unsaved changes.

About Creating a Device Alias

When you perform the first device alias task (regardless of which device alias task), the fabric is automatically locked for the device alias feature. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the effective database is obtained and used as the pending database. Modifications from this point on are made to the pending database. The pending database remains in effect until you commit the modifications to the pending database or discard (**abort**) the changes to the pending database.

Creating a Device Alias

To lock the fabric and create a device alias in the pending database using Fabric Manager, follow these steps:

Step 1 Expand **End Devices** and then select **Device Alias** in the Physical Attributes pane.

You see the device alias configuration in the Information pane.

Step 2 Click the **Configuration** tab and click the **Create Row** icon.

You see the Device Alias Creation dialog box in Figure 27-2.

Figure 27-2 Create Device Alias Dialog Box

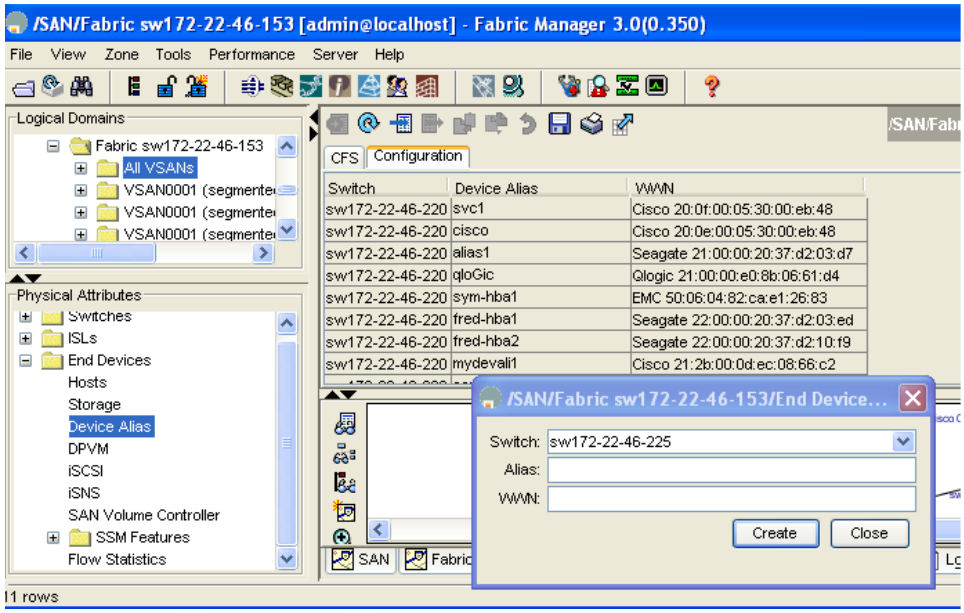

Step 3 Select a switch from the drop-down menu.

Step 4 Complete the Alias name and pWWN fields.

Step 5 Click **Create** to create this alias or click **Close** to discard any unsaved changes.

Committing Changes

If you commit the changes made to the pending database, the following events occur:

1. The pending database content overwrites the effective database content.
2. The pending database is emptied of its contents.
3. The fabric lock is released for this feature.

To commit the changes to the device alias database using Fabric Manager, follow these steps:

Step 1 Expand **End Devices** and then select **Device Alias** in the Physical Attributes pane.

You see the device alias configuration in the Information pane. The **CFS** tab is the default tab.

Step 2 Select **enable** from the Global drop-down menus to enabled switch aliases.

Step 3 Select **commit** from the Config Action drop-down menu for the newly enabled switches.

Step 4 Click **Apply Changes** to commit and distribute these changes or click **Undo Changes** to discard any unsaved changes.

Discarding Changes

If you discard changes made to the pending database, the following events occur:

1. The effective database contents remain unaffected.
2. The pending database is emptied of its contents.
3. The fabric lock is released for this feature.

To discard the device alias session using Fabric Manager, follow these steps:

Step 1 Expand **End Devices** and then select **Device Alias** in the Physical Attributes pane.

You see the device alias configuration in the Information pane. The **CFS** tab is the default tab.

Step 2 Select **abort** from the Config Action drop-down menu.

Step 3 Click **Apply Changes** to discard the session.

Legacy Zone Alias Conversion

You can import legacy zone alias configurations to use this feature without loosing data, if they satisfy the following restrictions:

- Each zone alias has only one member.
- The member type is pWWN.
- The name and definition of the zone alias should not be the same as any existing device alias name.

If any name conflict exists, the zone aliases are not imported.

Tip Ensure to copy any required zone aliases to the device alias database as required by your configuration.

When an import operation is complete, the modified alias database is distributed to all other switches in the physical fabric when you perform the **commit** operation. At this time if you do not want to distribute the configuration to other switches in the fabric, you can perform the **abort** operation and the merge changes are completely discarded.

This section includes the following topics:

Using Device Aliases or FC Aliases

You can change whether Fabric Manager uses FC aliases or global device aliases from Fabric Manager client without restarting Fabric Manager Server.

To change whether Fabric Manager uses FC aliases or global device aliases, follow these steps:

Step 1 Click **Server > Admin**.

You see the Admin dialog box in Figure 27-3.

Figure 27-3 Server Admin Dialog Box

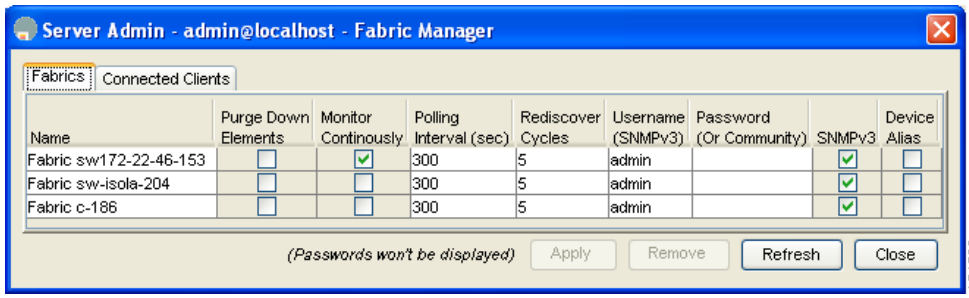

Step 2 For each fabric that you are monitoring with Fabric Manager Server, check the **Device Alias** check box to use global device aliases, or uncheck to use FC aliases.

Step 3 Click **Apply** to save these changes or click **Close** to exit the dialog box without saving any changes.

Device Alias Statistics Cleanup

To clear device alias statistics (for debugging purposes), refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

Database Merge Guidelines

Refer to the "CFS Merge Support" section on page 12-9 for detailed concepts.

When merging two device alias databases, follow these guidelines:

- Verify that two device aliases with different names are not mapped to the same pWWN.
- Verify that two identical pWWNs are not mapped to two different device aliases.
- Verify that the combined number of the device aliases in both databases does not exceed 8191 (8K). For example, if database N has 6000 device aliases and database M has 2192 device aliases, this merge operation will fail.

Default Settings

Table 27-2 lists the default settings for device alias parameters.

Table 27-2 Default Device Alias Parameters

Parameters	Default
Database in use	Effective database.
Database to accept changes	Pending database.
Device alias fabric lock state	Locked with the first device alias task.

CHAPTER 28

Configuring Fibre Channel Routing Services and Protocols

Fabric Shortest Path First (FSPF) is the standard path selection protocol used by Fibre Channel fabrics. The FSPF feature is enabled by default on all Fibre Channel switches. Except in configurations that require special consideration, you do not need to configure any FSPF services. FSPF automatically calculates the best path between any two switches in a fabric. Specifically, FSPF is used to:

- Dynamically compute routes throughout a fabric by establishing the shortest and quickest path between any two switches.
- Select an alternative path in the event of the failure of a given path. FSPF supports multiple paths and automatically computes an alternative path around a failed link. It provides a preferred route when two equal paths are available.

This chapter provides details on Fibre Channel routing services and protocols. It includes the following sections:

About FSPF

FSPF is the protocol currently standardized by the T11 committee for routing in Fibre Channel networks. The FSPF protocol has the following characteristics and features:

- Supports multipath routing.
- Bases path status on a link state protocol.
- Routes hop by hop, based only on the domain ID.
- Runs only on E ports or TE ports and provides a loop free topology.
- Runs on a per VSAN basis. Connectivity in a given VSAN in a fabric is guaranteed only for the switches configured in that VSAN.
- Uses a topology database to keep track of the state of the links on all switches in the fabric and associates a cost with each link.
- Guarantees a fast reconvergence time in case of a topology change. Uses the standard Dijkstra's algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra's algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.

FSPF Examples

This section provides examples of topologies and applications that demonstrate the benefits of FSPF.

Note The FSPF feature can be used on any topology.

Fault Tolerant Fabric

Figure 28-1 depicts a fault tolerant fabric using a partial mesh topology. If a link goes down anywhere in the fabric, any switch can still communicate with all others in the fabric. In the same way, if any switch goes down, the connectivity of the rest of the fabric is preserved.

Figure 28-1 Fault Tolerant Fabric

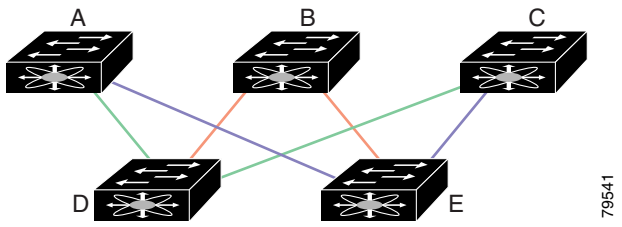

For example, if all links are of equal speed, the FSPF calculates two equal paths from A to C: A-D-C (green) and A-E-C (blue).

Redundant Links

To further improve on the topology in Figure 28-1, each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. Figure 28-2 shows this arrangement. Because switches in the Cisco MDS 9000 Family support PortChanneling, each pair of physical links can appear to the FSPF protocol as one single logical link.

By bundling pairs of physical links, FSPF efficiency is considerably improved by the reduced database size and the frequency of link updates. Once physical links are aggregated, failures are not attached to a single link but to the entire PortChannel. This configuration also improves the resiliency of the network. The failure of a link in a PortChannel does not trigger a route change, thereby reducing the risks of routing loops, traffic loss, or fabric downtime for route reconfiguration.

Figure 28-2 Fault Tolerant Fabric with Redundant Links

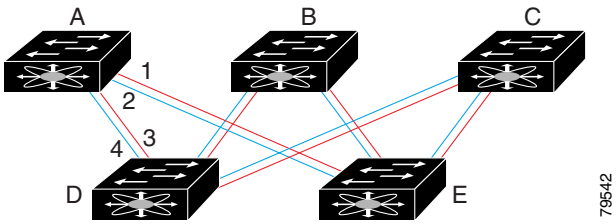

For example, if all links are of equal speed and no PortChannels exist, the FSPF calculates four equal paths from A to C: A1-E-C, A2-E-C, A3-D-C, and A4-D-C. If PortChannels exist, these paths are reduced to two.

Fail-Over Scenarios for PortChannels and FSPF Links

The SmartBits traffic generator was used to evaluate the scenarios displayed in Figure 28-3. Two links between switch 1 and switch 2 exist as either equal-cost ISLs or PortChannels. There is one flow from traffic generator 1 to traffic generator 2. The traffic was tested at 100% utilization of 1 Gbps in two scenarios:

- Disabling the traffic link by either physically removing the cable (see Table 28-1).
- Shutting down either switch 1 or switch 2 (see Table 28-2).

Figure 28-3 Fail-Over Scenario Using Traffic Generators

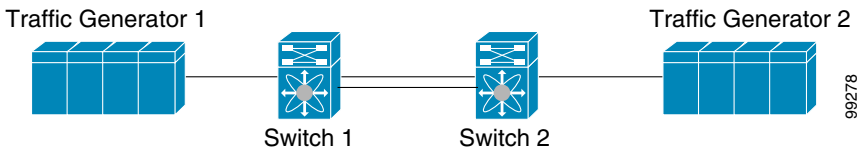

Table 28-1 Physically Removing the Cable for the SmartBits Scenario

PortChannel Scenario		FSPF Scenario (Equal cost ISL)	
Switch 1	**Switch 2**	**Switch 1**	**Switch 2**
110 msec (~2K frame drops)		130+ msec (~4k frame drops)	
100 msec (hold time when a signal loss is reported as mandated by the standard)			

Table 28-2 Shutting Down the Switch for the SmartBits Scenario

PortChannel Scenario		FSPF Scenario (Equal cost ISL)	
Switch 1	**Switch 2**	**Switch 1**	**Switch 2**
~0 msec (~8 frame drops)	110 msec (~2K frame drops)	130+ msec (~4K frame drops)	
No hold time needed	Signal loss on switch 1	No hold time needed	Signal loss on switch 1

FSPF Global Configuration

By default, FSPF is enabled on switches in the Cisco MDS 9000 Family.

Some FSPF features can be globally configured in each VSAN. By configuring a feature for the entire VSAN, you do not have to specify the VSAN number for every command. This global configuration feature also reduces the chance of typing errors or other minor configuration errors.

Note FSPF is enabled by default. Generally, you do not need to configure these advanced features.

Caution The default for the backbone region is 0 (zero). You do not need to change this setting unless your region is different from the default. If you are operating with other vendors using the backbone region, you can change this default to be compatible with those settings.

This section includes the following topics:

About SPF Computational Hold Times

The SPF computational hold time sets the minimum time between two consecutive SPF computations on the VSAN. Setting this to a small value means that FSPF reacts faster to any fabric changes by recomputing paths on the VSAN. A small SPF computational hold time uses more switch CPU time.

About Link State Records

Each time a new switch enters the fabric, a link state record (LSR) is sent to the neighboring switches, and then flooded throughout the fabric. Table 28-3 displays the default settings for switch responses.

Table 28-3 LSR Default Settings

LSR Option	Default	Description
Acknowledgment interval (RxmtInterval)	5 seconds	The time a switch waits for an acknowledgment from the LSR before retransmission.
Refresh time (LSRefreshTime)	30 minutes	The time a switch waits before sending an LSR refresh transmission.
Maximum age (MaxAge)	60 minutes	The time a switch waits before dropping the LSR from the database.

The LSR minimum arrival time is the period between receiving LSR updates on this VSAN. Any LSR updates that arrive before the LSR minimum arrival time are discarded.

The LSR minimum interval time is the frequency at which this switch sends LSR updates on a VSAN.

Configuring FSPF on a VSAN

To configure an FSPF feature for the entire VSAN using Fabric Manager, follow these steps:

Step 1 Expand a Fabric, expand a VSAN and select **FSPF** for a VSAN that you want to configure for FSPF.

You see the FSPF configuration in the Information pane in Figure 28-4.

Figure 28-4 FSPF General Information

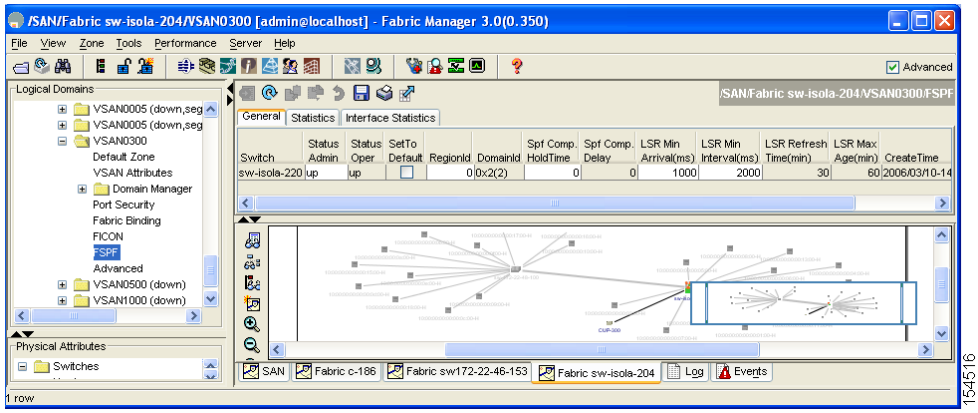

Step 2 The RegionID, Spf Comp Holdtime, LSR Min Arrival, and LSR Min Interval field values are applied across all interfaces on the VSAN. You can change them here or, if they don't exist, create them here.

Step 3 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

Resetting FSPF to the Default Configuration

To return the FSPF VSAN global configuration to its factory default using Fabric Manager, follow these steps:

Step 1 Expand a Fabric, expand a VSAN and select **FSPF** for a VSAN that you want to configure for FSPF.

You see the FSPF configuration in the Information pane in Figure 28-5.

Figure 28-5 FSPF General Information

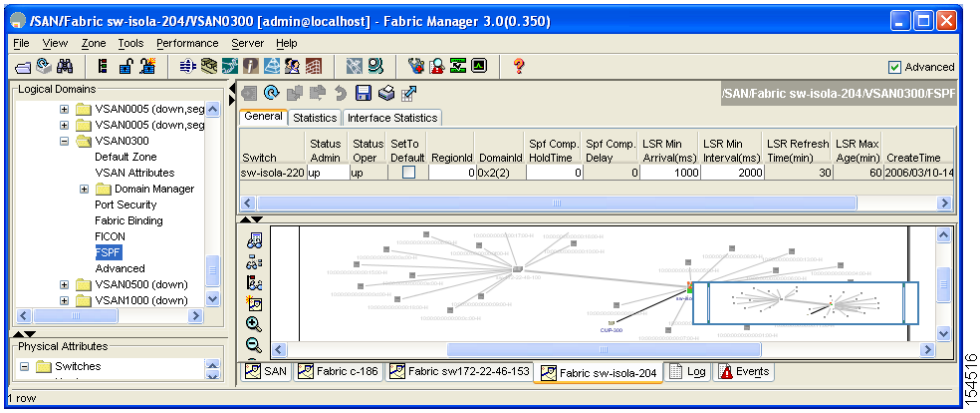

Step 2 Check the **SetToDefault** check box for a switch (see Figure 28-5).

Step 3 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

Enabling or Disabling FSPF

To enable or disable FSPF using Fabric Manager, follow these steps:

Step 1 Expand a Fabric, expand a VSAN and select **FSPF** for a VSAN that you want to configure for FSPF.

You see the FSPF configuration in the Information pane in Figure 28-6.

Figure 28-6 FSPF General Information

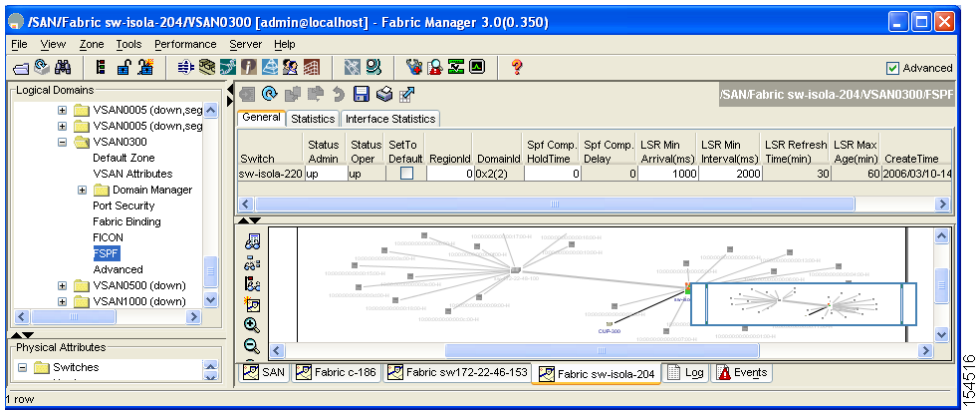

Step 2 Set the Status Admin drop-down menu to **up** to enable FSPF or to **down** to disable FSPF (see Figure 28-6).

Step 3 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

FSPF Interface Configuration

Several FSPF commands are available on a per interface basis. These configuration procedures apply to an interface in a specific VSAN.

This section includes the following topics:

About FSPF Link Cost

FSPF tracks the state of links on all switches in the fabric, associates a cost with each link in its database, and then chooses the path with a minimal cost. The cost associated with an interface can be administratively changed to implement the FSPF route selection. The integer value to specify cost can range from 1 to 65,535. The default cost for 1 Gbps is 1000 and for 2 Gbps is 500.

Configuring FSPF Link Cost

To configure FSPF link cost using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **Interfaces** and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **FSPF** tab.

You see the FSPF interface configuration in the Information pane in Figure 28-7.

Figure 28-7 Fibre Channel Physical FSPF Interface

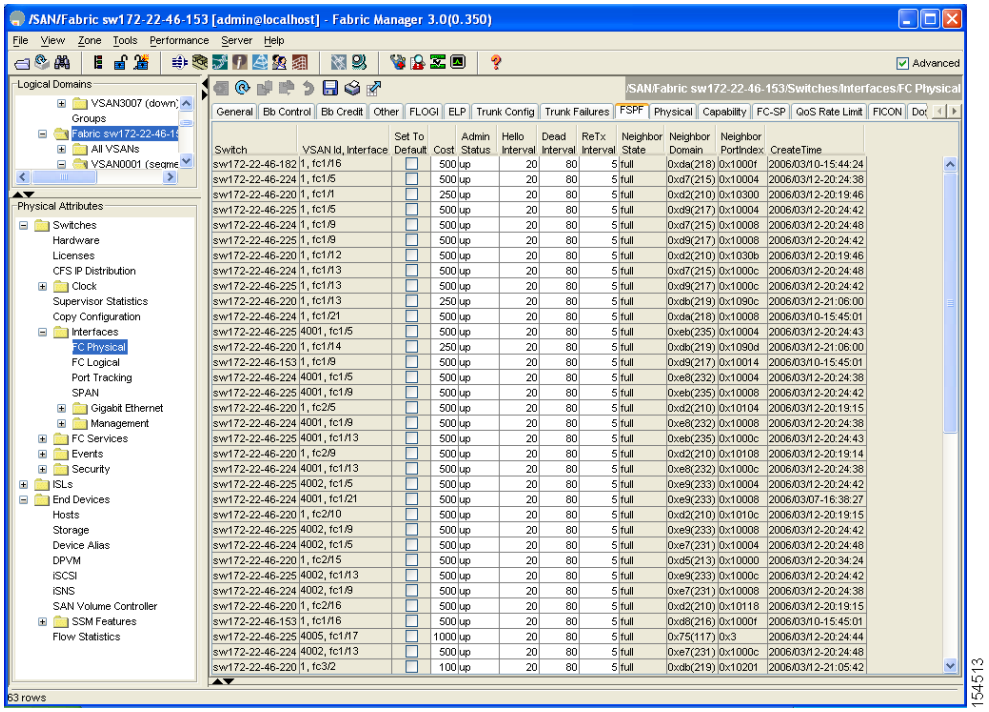

Step 3 Double-click in the Cost field of a switch and change the value (see Figure 28-7).

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

About Hello Time Intervals

You can set the FSPF Hello time interval to specify the interval between the periodic hello messages sent to verify the health of the link. The integer value can range from 1 to 65,535 seconds.

Note This value must be the same in the ports at both ends of the ISL.

Configuring Hello Time Intervals

To configure the FSPF Hello time interval using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **Interfaces** and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **FSPF** tab.

You see the FSPF interface configuration in the Information pane (see Figure 28-8).

Figure 28-8 Fibre Channel Physical FSPF Interface

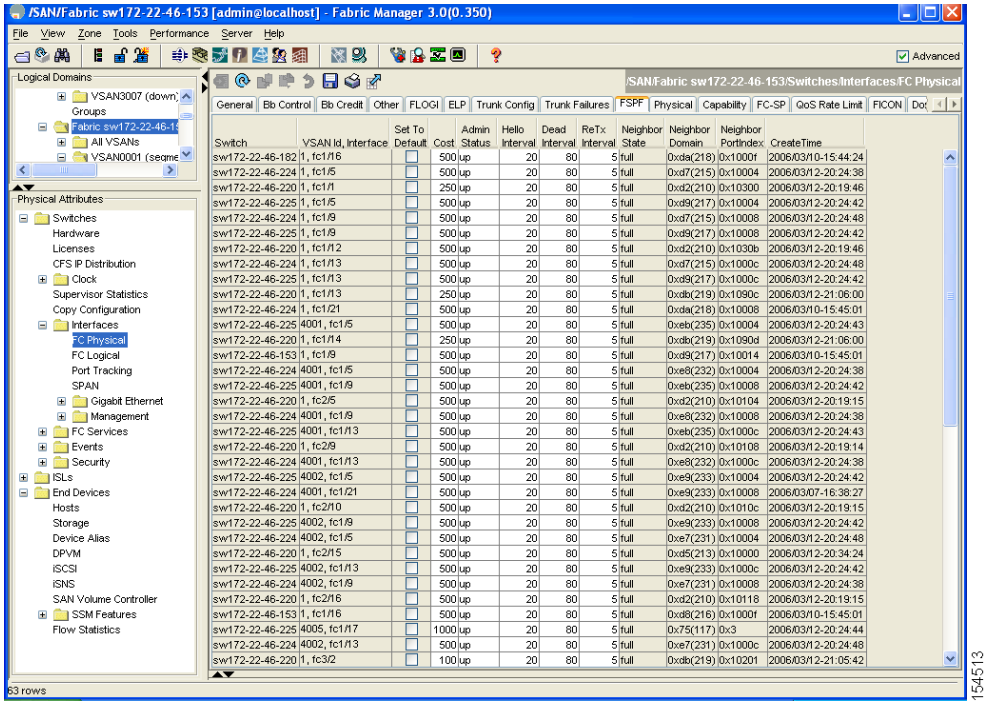

Step 3 Change the Hello Interval field (see Figure 28-8) for a switch.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

About Dead Time Intervals

You can set the FSPF dead time interval to specify the maximum interval for which a hello message must be received before the neighbor is considered lost and removed from the database. The integer value can range from 1 to 65,535 seconds.

Note This value must be the same in the ports at both ends of the ISL.

Caution An error is reported at the command prompt if the configured dead time interval is less than the hello time interval.

Configuring Dead Time Intervals

To configure the FSPF dead time interval using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **Interfaces** and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **FSPF** tab.

You see the FSPF interface configuration in the Information pane (see Figure 28-9).

Figure 28-9 Fibre Channel Physical FSPF Interface

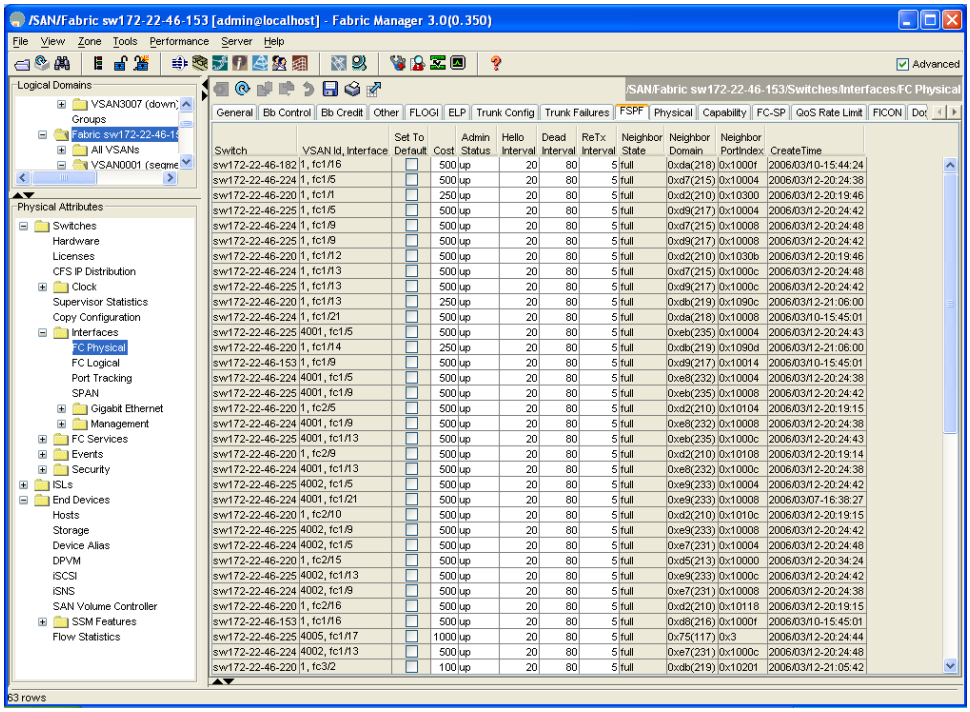

Step 3 Double-click the Dead Interval field for a switch and provide a new value (see Figure 28-9).

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

About Retransmitting Intervals

You can specify the time after which an unacknowledged link state update should be transmitted on the interface. The integer value to specify retransmit intervals can range from 1 to 65,535 seconds.

Note This value must be the same on the switches on both ends of the interface.

Configuring Retransmitting Intervals

To configure the FSPF retransmit time interval using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **Interfaces**, and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **FSPF** tab.

You see the FSPF interface configuration in the Information pane shown in Figure 28-10.

Figure 28-10 Fibre Channel Physical FSPF Interface

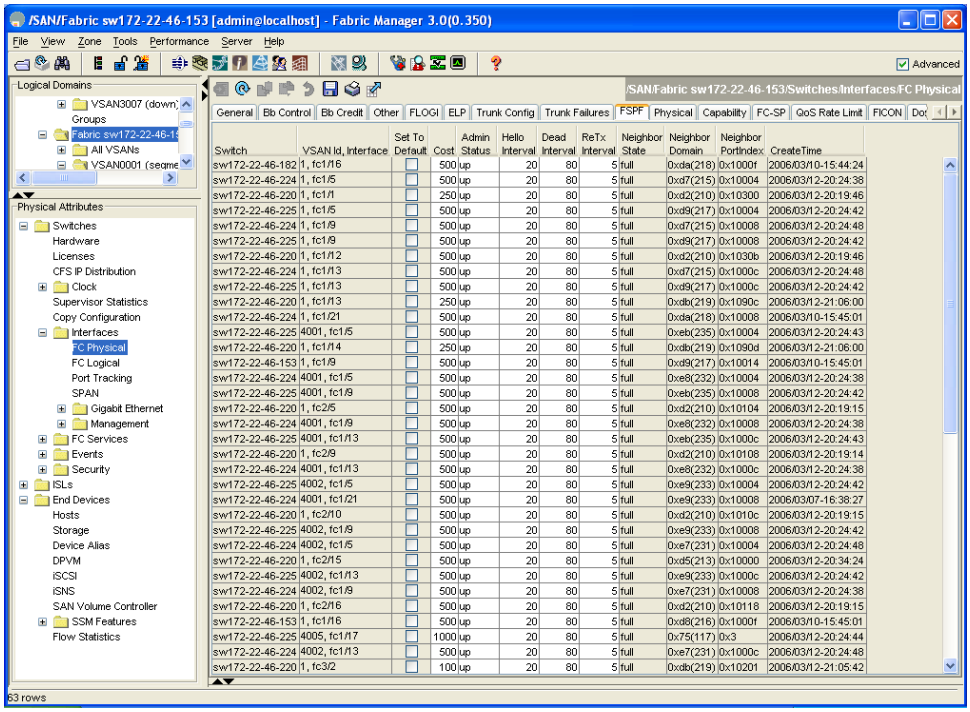

Step 3 Double-click the ReTx Interval field and enter a value.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

About Disabling FSPF for Specific Interfaces

You can disable the FSPF protocol for selected interfaces. By default, FSPF is enabled on all E ports and TE ports. This default can be disabled by setting the interface as passive.

Note FSPF must be enabled at both ends of the interface for the protocol to work.

Disabling FSPF for Specific Interfaces

To disable FSPF for a specific interface using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **Interfaces** and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **FSPF** tab.

You see the FSPF interface configuration in the Information pane shown in Figure 28-11.

Figure 28-11 Fibre Channel Physical FSPF Interface

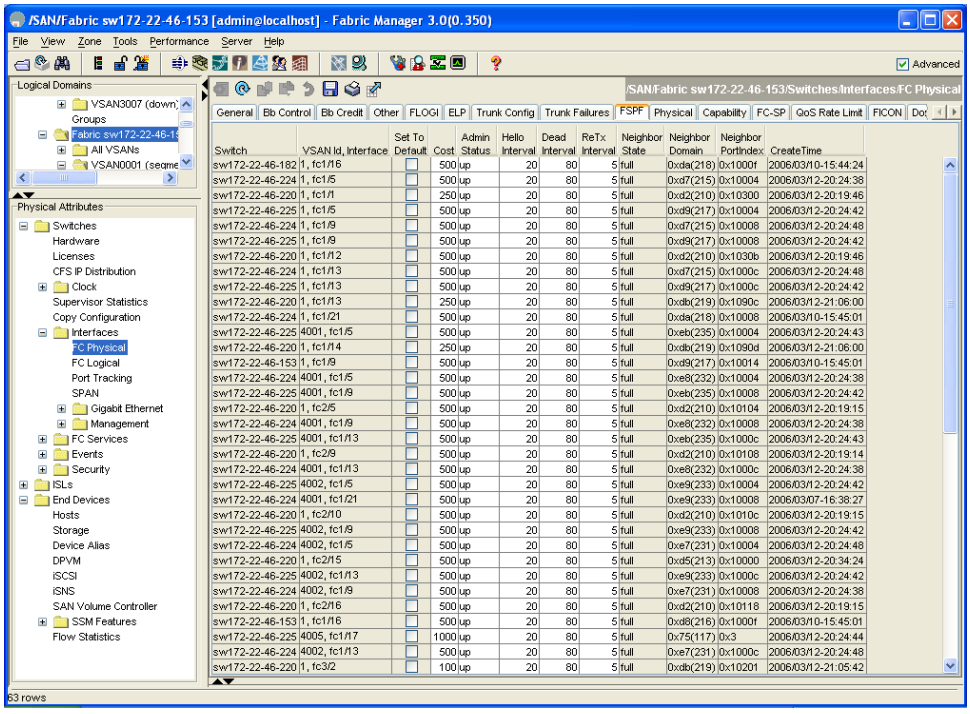

Step 3 Set a switch Admin Status drop-down menu to **down**.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

Displaying the FSPF Database

The FSPF database for a specified VSAN includes the following information:

- Link State Record (LSR) type
- Domain ID of the LSR owner
- Domain ID of the advertising router
- LSR age
- LSR incarnation member
- Number of links

To display the FSPF database using Device Manager, follow these steps:

Step 1 Choose **FC > Advanced > FSPF**.

You see the FSPF dialog box shown in Figure 28-12.

Figure 28-12 FSPF Configuration Dialog Box in Device Manager

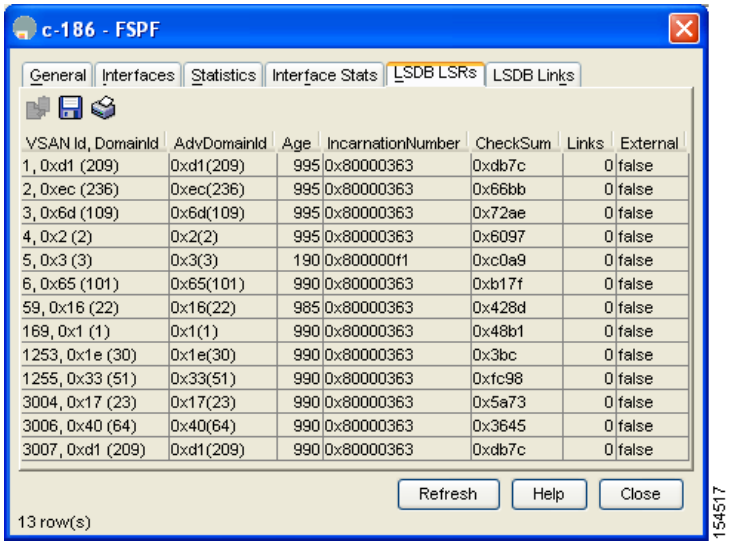

Step 2 Click the **LSDB LSRs** tab.

You see the FSPF database information shown in Figure 28-13.

Figure 28-13 FSPF Database Information in the LSDB LSRs Tab

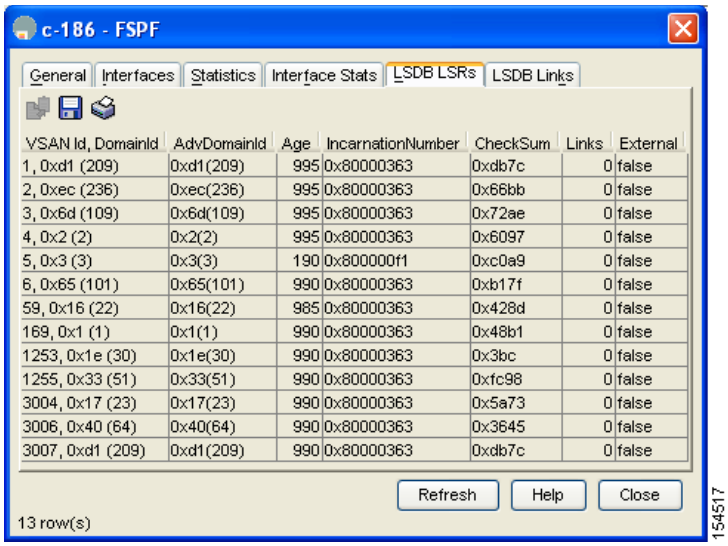

Viewing FSPF Statistics

To view FSPF statistics using Fabric Manager, follow these steps:

Step 1 Expand a Fabric, expand a VSAN, and then select **FSPF** in the Logical Domains pane.

You see the FSPF configuration dialog box.

Step 2 Click the **Statistics** tab.

You see the FSPF VSAN statistics in the Information pane (see Figure 28-14).

Figure 28-14 FSPF VSAN Statistics

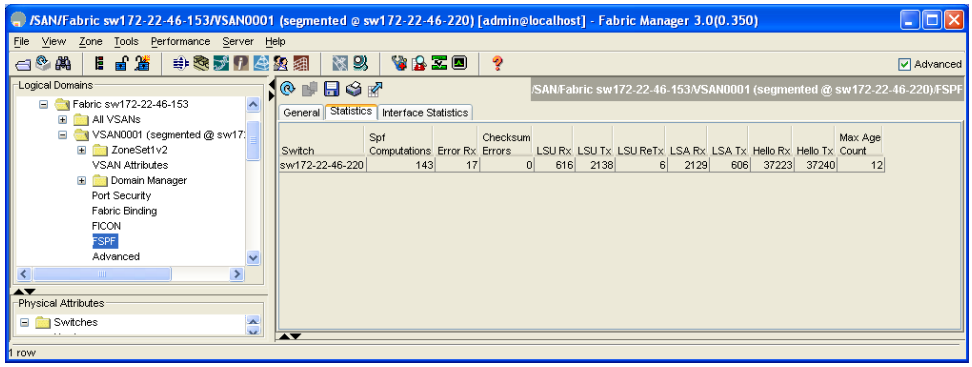

Step 3 Click the **Interface Statistics** tab.

You see the FSPF interface statistics in the Information pane.

FSPF Routes

FSPF routes traffic across the fabric, based on entries in the FSPF database. These routes can be learned dynamically, or configured statically.

This section includes the following topics:

About Fibre Channel Routes

Each port implements forwarding logic, which forwards frames based on its FC ID. To configure the FC ID for the specified interface and domain, you can configure the specified route (for example FC ID 111211 and domain ID 3) in the switch with domain ID 1 (see Figure 28-15).

Figure 28-15 Fibre Channel Routes

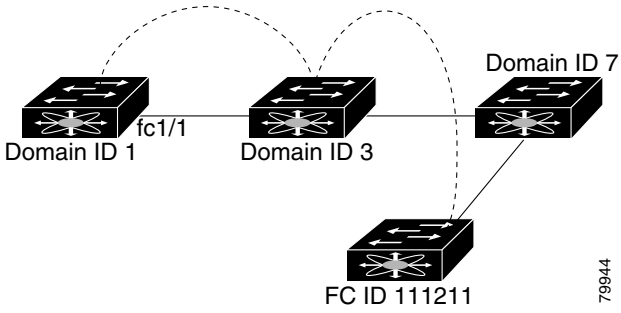

Note Other than in VSANs, runtime checks are not performed on configured and suspended static routes.

Configuring Fibre Channel Routes

To configure a Fibre Channel route using Device Manager, follow these steps:

Step 1 Click **FC > Advanced > Routes**.

You see the FC Static Route Configuration dialog box shown in Figure 28-16.

Figure 28-16 Fibre Channel Static Route Configuration

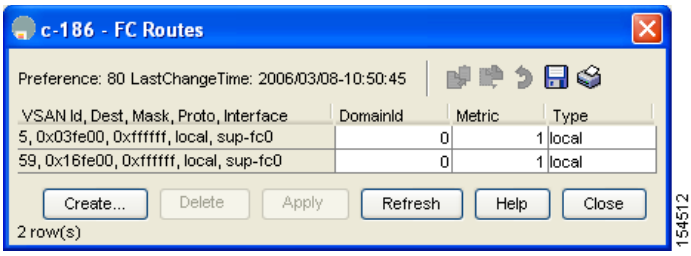

Step 2 Click **Create** to create a static route.

You see the Create Route dialog box shown in Figure 28-17.

Figure 28-17 Create Fibre Channel Route

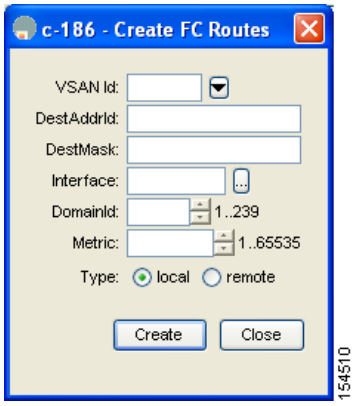

Step 3 Select the VSAN ID that you are configuring this route on (see Figure 28-17).

Step 4 Fill in the destination address and destination mask for the device you are configuring a route to (see Figure 28-17).

Step 5 Select the interface that you want to use to reach this destination (see Figure 28-17).

Step 6 Select the next hop domain ID and route metric (see Figure 28-17).

Step 7 Select either the **local** or **remote** radio button (see Figure 28-17).

Step 8 Click **Create** to save these changes or click **Close** to discard any unsaved changes (see Figure 28-17).

About Broadcast and Multicast Routing

Broadcast and multicast in a Fibre Channel fabric uses the concept of a distribution tree to reach all switches in the fabric.

FSPF provides the topology information to compute the distribution tree. Fibre Channel defines 256 multicast groups and one broadcast address for each VSAN. Switches in the Cisco MDS 9000 Family only use broadcast routing. By default, they use the principal switch as the root node to derive a loop-free distribution tree for multicast and broadcast routing in a VSAN.

Caution All switches in the fabric should run the same multicast and broadcast distribution tree algorithm to ensure the same distribution tree.

To interoperate with other vendor switches (following FC-SW3 guidelines), the SAN-OS software uses the lowest domain switch as the root to compute the multicast tree in interop mode.

About Multicast Root Switch

By default, the native (non-interop) mode uses the principal switch as the root. If you change the default, be sure to configure the same mode in all switches in the fabric. Otherwise, multicast traffic could face potential loop and frame-drop problems.

Note The operational mode can be different from the configured interop mode. The interop mode always uses the lowest domain switch as the root.

Setting the Multicast Root Switch

To use the lowest domain switch for the multicast tree computation using Fabric Manager, follow these steps:

Step 1 Expand a fabric, expand a VSAN, and then select **Advanced** for the VSAN that you want to configure FSPF on.

You see the advanced Fibre Channel configuration in the Information pane.

Step 2 Select the **Multicast Root** tab.

You see the multicast root configuration in the Information pane shown in Figure 28-18.

Figure 28-18 Multicast Root Configuration

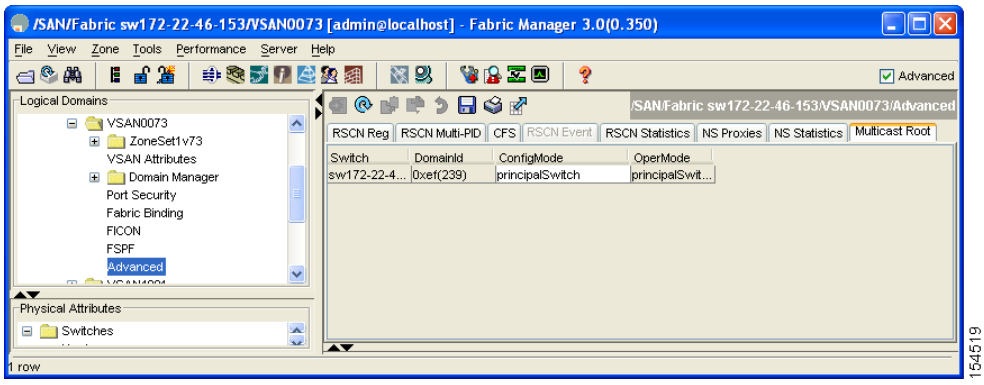

Step 3 Set the Config Mode drop-down menu to **lowestDomainSwitch** (see Figure 28-18).

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

In-Order Delivery

In-Order Delivery (IOD) of data frames guarantees frame delivery to a destination in the same order that they were sent by the originator.

Some Fibre Channel protocols or applications cannot handle out-of-order frame delivery. In these cases, switches in the Cisco MDS 9000 Family preserve frame ordering in the frame flow. The source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) identify the flow of the frame.

On any given switch with IOD enabled, all frames received by a specific ingress port and destined to a certain egress port are always delivered in the same order in which they were received.

Use IOD only if your environment cannot support out-of-order frame delivery.

Tip If you enable the in-order delivery feature, the graceful shutdown feature is not implemented.

This section includes the following topics:

About Reordering Network Frames

When you experience a route change in the network, the new selected path may be faster or less congested than the old route (see Figure 28-19).

Figure 28-19 Route Change Delivery

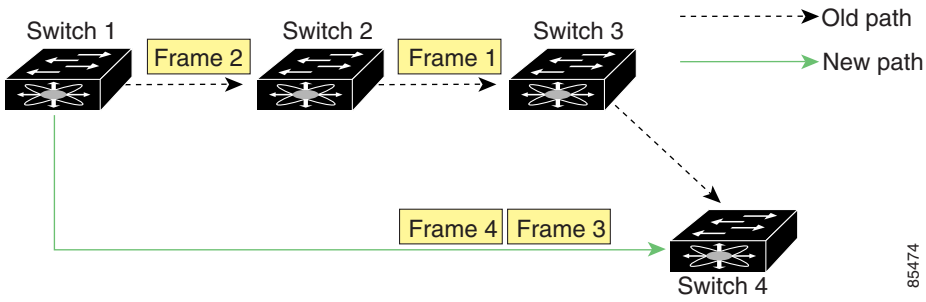

In Figure 28-19, the new path from Switch 1 to Switch 4 is faster. In this scenario, Frame 3 and Frame 4 may be delivered before Frame 1 and Frame 2.

If IOD is enabled, the frames within the network are treated as follows:

- Frames in the network are delivered in the order in which they are transmitted.
- Frames that cannot be delivered in order within the network latency drop period are dropped inside the network.

About Reordering PortChannel Frames

When a link change occurs in a PortChannel, the frames for the same exchange or the same flow can switch from one path to another faster path (see Figure 28-20).

Figure 28-20 Link Congestion Delivery

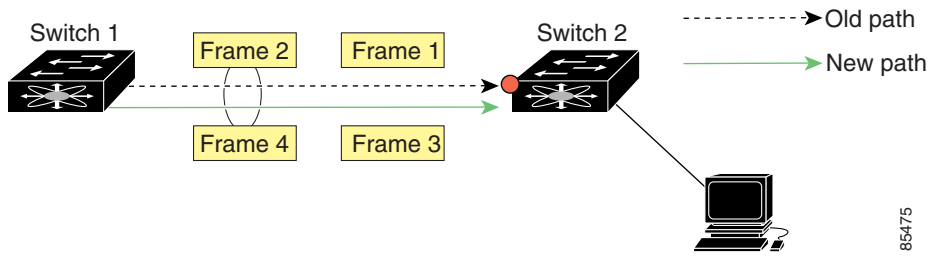

In Figure 28-20, the port of the old path (red dot) is congested. In this scenario, Frame 3 and Frame 4 can be delivered before Frame 1 and Frame 2.

As of Cisco SAN-OS Release 3.0(1), the IOD feature attempts to minimize the number of frames dropped during PortChannel link changes when IOD is enabled by sending a request to the remote switch on the PortChannel to flush all frames for this PortChannel.

Note Both switches on the PortChannel must be running Cisco SAN-OS Release 3.0(1) for this IOD enhancement. For earlier releases, IOD waits for the switch latency period before sending new frames.

When IOD is enabled and a PortChannel link change occurs, the frames crossing the PortChannel are treated as follows:

- Frames using the old path are delivered before new frames are accepted.
- The new frames are delivered through the new path after all old frames are flushed.

Frames that cannot be delivered in order through the old path within the switch latency drop period are dropped. See the "Configuring the Drop Latency Time" section on page 28-21.

About Enabling In-Order Delivery

You can enable IOD for a specific VSAN or for the entire switch. By default, IOD is disabled on switches in the Cisco MDS 9000 Family.

Tip We recommend that you only enable this feature when devices that cannot handle any out-of-order frames are present in the switch. Load-balancing algorithms within the Cisco MDS 9000 Family ensure that frames are delivered in order during normal fabric operation. The load-balancing algorithms based

on source FC ID, destination FC ID, and exchange ID are enforced in hardware without any performance degradation. However, if the fabric encounters a failure and this feature is enabled, the recovery will be delayed because of an intentional pausing of fabric forwarding to purge the fabric of resident frames that could potentially be forwarded out-of-order.

Enabling IOD Globally

Only enable IOD globally if this is a requirement across your entire fabric. Otherwise, enable IOD only for the VSANs that require this feature.

Note Enable IOD on the entire switch before performing a downgrade to Cisco MDS SAN-OS Release1.3(3) or earlier.

To enable IOD for the switch using Fabric Manager, follow these steps:

Step 1 Expand a fabric and select **All VSANS**.

You see the VSAN configuration in the Information pane.

Step 2 Click the **Attributes** tab.

You see the general VSAN attributes in the Information pane shown in Figure 28-21.

Figure 28-21 General VSAN Attributes

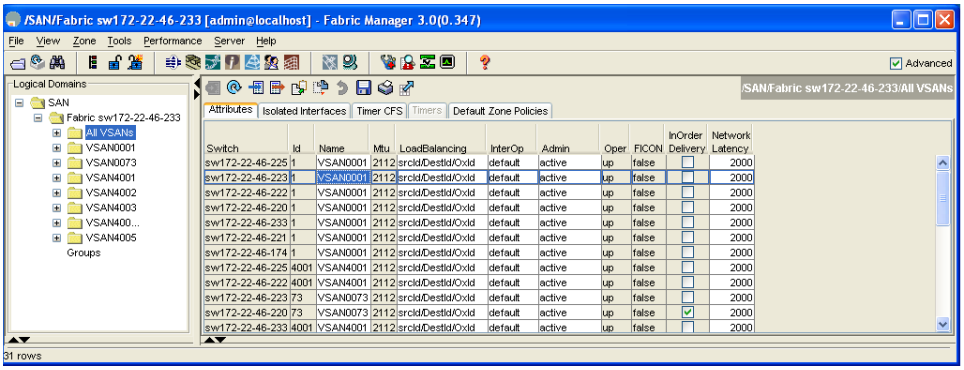

Step 3 Check the **InOrder Delivery** check box to enable IOD for each switch (see Figure 28-21).

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

Enabling IOD for a VSAN

When you create a VSAN, that VSAN automatically inherits the global IOD value. You can override this global value by enabling or disabling IOD for the new VSAN.

To enable IOD for a specific VSAN, follow these steps:

Step 1 Expand a fabric and select **All VSANS**.

You see the VSAN configuration in the Information pane.

Step 2 Select the **Attributes** tab.

You see the general VSAN attributes in the Information pane shown in Figure 28-22.

Figure 28-22 General VSAN Attributes

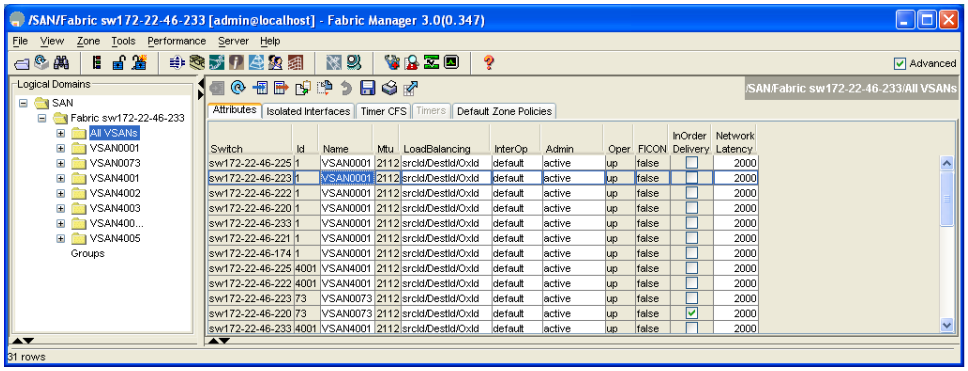

Step 3 Check the **InOrder Delivery** check box to enable IOD for the switch (see Figure 28-22).

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

Configuring the Drop Latency Time

You can change the default latency time for either the entire switch or a specified VSAN in a switch.

To configure the drop latency time for a switch using Fabric Manager, follow these steps:

Step 1 Expand a fabric and select **All VSANS**.

You see the VSAN configuration in the Information pane.

Step 2 Select the **Attributes** tab.

You see the general VSAN attributes in the Information pane shown in Figure 28-23.

Figure 28-23 General VSAN Attributes

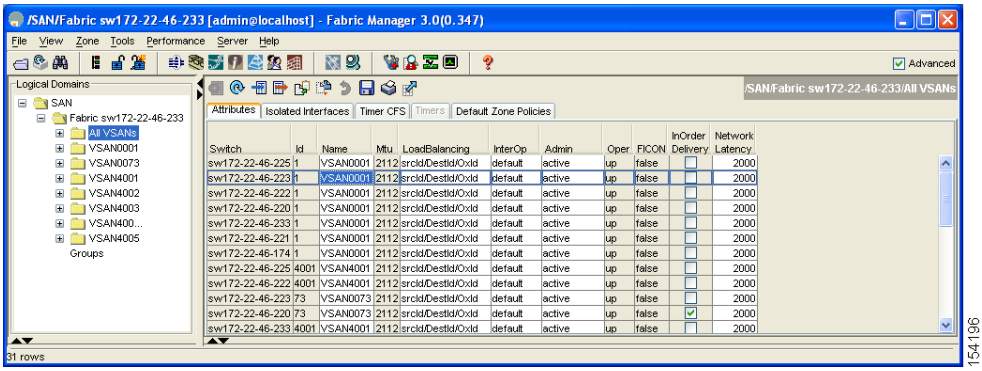

Step 3 Double-click the Network Latency field and change the value (see Figure 28-23).

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

Flow Statistics Configuration

Flow statistics count the ingress traffic in the aggregated statistics table. You can collect two kinds of statistics:

- Aggregated flow statistics to count the traffic for a VSAN.
- Flow statistics to count the traffic for a source and destination ID pair in a VSAN.

This section includes the following topics:

About Flow Statistics

If you enable flow counters, you can enable a maximum of 1K entries for aggregate flow and flow statistics. Be sure to assign an unused flow index to a module for each new flow. Flow indexes can be repeated across modules. The number space for flow index is shared between the aggregate flow statistics and the flow statistics.

Counting Aggregate Flow Statistics

To count the aggregated flow statistics for a VSAN using Device Manager, follow these steps:

Step 1 Click **FC > Advanced > Flow Statistics**.

You see the Flow Statistics dialog box shown in Figure 28-24.

Figure 28-24 Flow Statistics dialog box

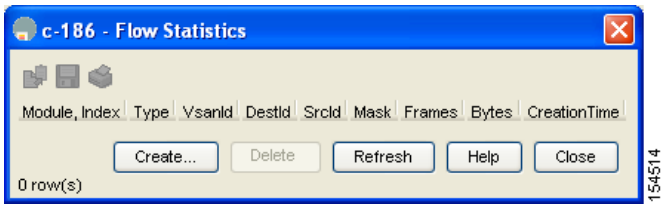

Step 2 Click **Create** to create a flow (see Figure 28-24).

You see the Create Flow dialog box shown in Figure 28-25.

Figure 28-25 Create Flow dialog box

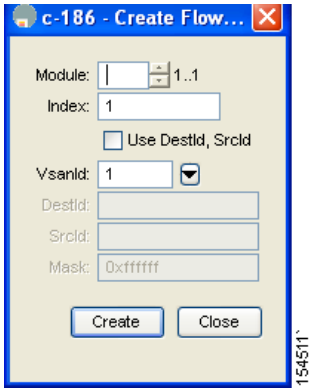

Step 3 Select the module, index, and VSAN ID for the flow you want to create (see Figure 28-25).

Step 4 Click **Create** to create this flow or click **Close** to discard any unsaved changes (see Figure 28-25).

Counting Individual Flow Statistics

To count the flow statistics for a source and destination FC ID in a VSAN using Fabric Manager, follow these steps:

Step 1 Click **FC > Advanced > Flow Statistics**.

You see the Flow Statistics dialog box in Figure 28-26.

Figure 28-26 Flow Statistics dialog box

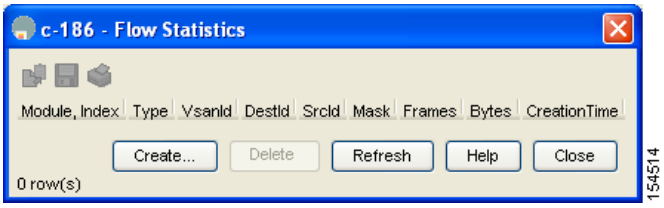

Step 2 Click **Create** to create a flow (see Figure 28-26).

You see the Create Flow dialog box shown in Figure 28-27.

Figure 28-27 Create Flow dialog box

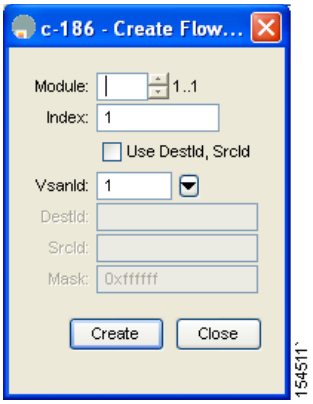

Step 3 Check the **Use DstId,SrcId** check box (see Figure 28-27).

Step 4 Complete the destination FC ID, source FC ID, and mask fields (see Figure 28-27).

Step 5 Click **Create** to create this flow or click **Close** to discard any unsaved changes (see Figure 28-27).

Default Settings

Table 28-4 lists the default settings for FSPF features.

Table 28-4 Default FSPF Settings

Parameters	Default
FSPF	Enabled on all E ports and TE ports.
SPF computation	Dynamic.
SPF hold time	0.
Backbone region	0.
Acknowledgment interval (RxmtInterval)	5 seconds.
Refresh time (LSRefreshTime)	30 minutes.
Maximum age (MaxAge)	60 minutes.
Hello interval	20 seconds.
Dead interval	80 seconds.
Distribution tree information	Derived from the principal switch (root node).
Routing table	FSPF stores up to 16 equal cost paths to a given destination.
Load balancing	Based on destination ID and source ID on different, equal cost paths.
In-order delivery	Disabled.
Drop latency	Disabled.
Static route cost	If the cost (metric) of the route is not specified, the default is 10.
Remote destination switch	If the remote destination switch is not specified, the default is direct.
Multicast routing	Uses the principal switch to compute the multicast tree.

CHAPTER 29

Managing FLOGI, Name Server, FDMI, and RSCN Databases

This chapter describes the fabric login database, the name server features, the Fabric-Device Management Interface, and Registered State Change Notification (RSCN) information provided in the Cisco MDS 9000 Family. It includes the following sections:

FLOGI

In a Fibre Channel fabric, each host or disk requires an FC ID. If the required device is displayed in the FLOGI table, the fabric login is successful. Examine the FLOGI database on a switch that is directly connected to the host HBA and connected ports. See the "Default Company ID list" section on page 32-16 and the "Switch Interoperability" section on page 32-17.

Displaying FLOGI Details

To verify that a storage device is in the fabric login (FLOGI) table using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **Interfaces** and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **FLOGI** tab.

You see all end devices that are logged into the fabric (see Figure 29-1).

Figure 29-1 FLOGI Physical Interfaces

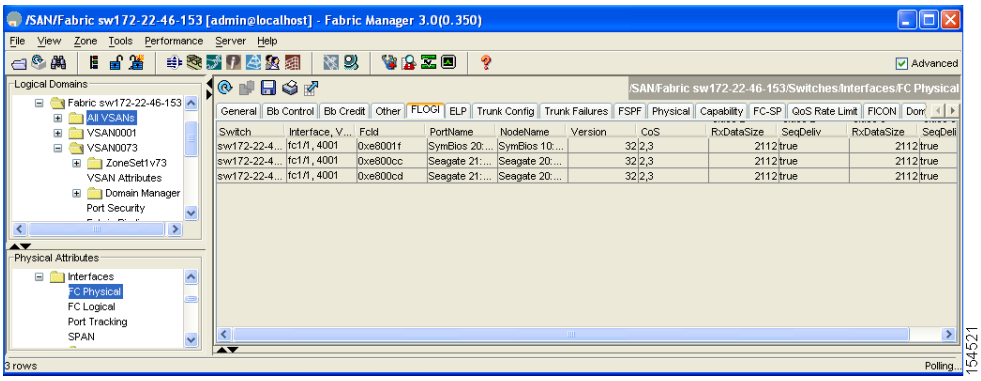

Name Server Proxy

The name server functionality maintains a database containing the attributes for all hosts and storage devices in each VSAN. Name servers allow a database entry to be modified by a device that originally registered the information.

The proxy feature is useful when you wish to modify (update or delete) the contents of a database entry that was previously registered by a different device.

This section includes the following topics:

About Registering Name Server Proxies

All name server registration requests come from the same port whose parameter is registered or changed. If it does not, then the request is rejected.

This authorization enables WWNs to register specific parameters for another node.

Registering Name Server Proxies

To register the name server proxy using Fabric Manager, follow these steps:

Step 1 Expand a fabric, expand a VSAN, and then select **Advanced**.

You see the VSAN advanced configuration in the Information pane.

Step 2 Click the **NS Proxies** tab.

You see the existing name server proxy for the selected VSAN in Figure 29-2.

Figure 29-2 Name Server Proxies

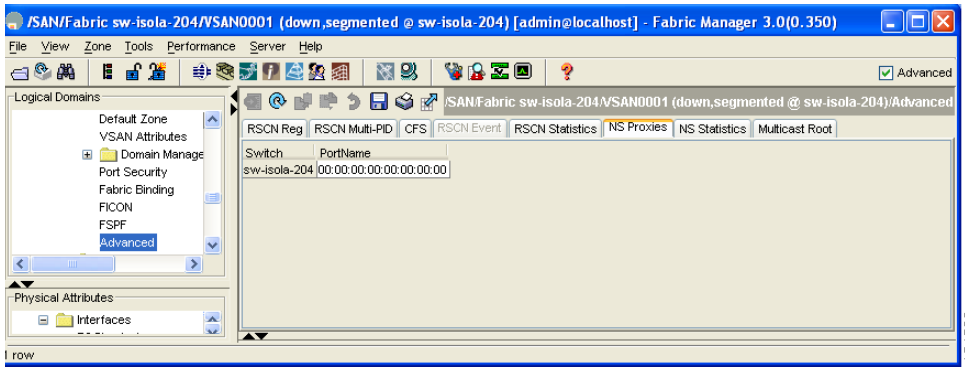

Step 3 Double-click the PortName field to register a new name server proxy (see Figure 29-2).

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to cancel any unsaved changes.

About Duplicate pWWN

You can prevent malicious or accidental log in using another device's pWWN. These pWWNs are allowed to log in to the fabric and replace the first device in the name server database.

Rejecting Duplicate pWWNs

To reject duplicate pWWNs, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

About Name Server Database Entries

The name server stores name entries for all hosts in the FCNS database. The name server permits an Nx port to register attributes during a PLOGI (to the name server) to obtain attributes of other hosts. These attributes are deregistered when the Nx port logs out either explicitly or implicitly.

In a multiswitch fabric configuration, the name server instances running on each switch shares information in a distributed database. One instance of the name server process runs on each switch.

Viewing Name Server Database Entries

To view the name server database using Device Manager, follow these steps:

Step 1 Click **FC > Name Server**.

You see the Name Server information in the Information pane in Figure 29-3.

Figure 29-3 Name Server Database Information

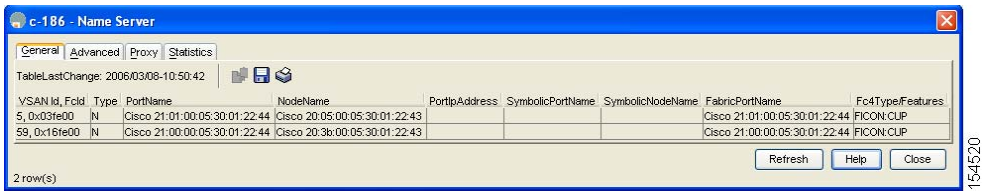

The General tab is the default tab; you see the name server database (see Figure 29-3).

Step 2 Click the **Statistics** tab.

You see the name server statistics.

FDMI

Cisco MDS 9000 Family switches provide support for the Fabric-Device Management Interface (FDMI) functionality, as described in the FC-GS-4 standard. FDMI enables management of devices such as Fibre Channel Host Bus Adapters (HBAs) through in-band communications. This addition complements the existing Fibre Channel name server and management server functions.

Using the FDMI functionality, the SAN-OS software can extract the following management information about attached HBAs and host operating systems without installing proprietary host agents:

- Manufacturer, model, and serial number
- Node name and node symbolic name
- Hardware, driver, and firmware versions
- Host operating system (OS) name and version number

All FDMI entries are stored in persistent storage and are retrieved when the FDMI process is started.

Displaying FDMI

To display the FDMI database information using Device Manager, choose **FC > Advanced > FDMI**. You see the FDMI dialog box.

RSCN

The Registered State Change Notification (RSCN) is a Fibre Channel service that informs hosts about changes in the fabric. Hosts can receive this information by registering with the fabric controller (through SCR). These notifications provide a timely indication of one or more of the following events:

- Disks joining or leaving the fabric.
- A name server registration change.
- A new zone enforcement.

- IP address change.
- Any other similar event that affects the operation of the host.

This section includes the following topics:

About RSCN Information

Apart from sending these events to registered hosts, a switch RSCN (SW-RSCN) is sent to all reachable switches in the fabric.

Note The switch sends an RSCN to notify registered nodes that a change has occurred. It is up to the nodes to query the name server again to obtain the new information. The details of the changed information are not delivered by the switch in the RSCN sent to the nodes.

Displaying RSCN Information

To display RSCN information using Fabric Manager, follow these steps:

Step 1 Expand a fabric, expand a VSAN and then select **Advanced**.

You see the VSAN advanced configuration in the Information pane.

Step 2 Select the **RSCN Reg** tab or the **RSCN Statistics** tab (see Figure 29-4).

Figure 29-4 RSCN Statistics

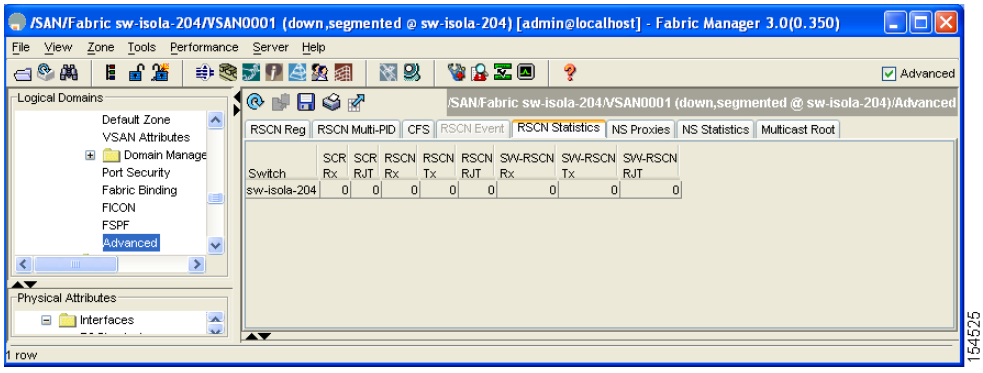

About the multi-pid Option

If the RSCN **multi-pid** option is enabled, then RSCNs generated to the registered Nx ports may contain more than one affected port IDs. In this case, zoning rules are applied before putting the multiple affected port IDs together in a single RSCN. By enabling this option, you can reduce the number of RSCNs. For example: Suppose you have two disks (D1, D2) and a host (H) connected to switch 1. Host H is registered to receive RSCNs. D1, D2 and H belong to the same zone. If disks D1 and D2 are online at the same time, then one of the following applies:

- The **multi-pid** option is disabled on switch 1: two RSCNs are generated to host H—one for the disk D1 and another for disk D2.
- The **multi-pid** format is enabled on switch 1: a single RSCN is generated to host H, and the RSCN payload lists the affected port IDs (in this case, both D1 and D2).

Note Some Nx ports may not understand multi-pid RSCN payloads. If so, disable the **multi-pid** RSCN option.

Configuring the multi-pid Option

To configure the multi-pid option using Fabric Manager, follow these steps:

Step 1 Expand a fabric, expand a VSAN and then select **Advanced**.

You see the VSAN advanced configuration in the Information pane.

Step 2 Click the **RSCN Multi-PID** tab.

You see the screen shown in Figure 29-5.

Figure 29-5 RSCN Multi-PID

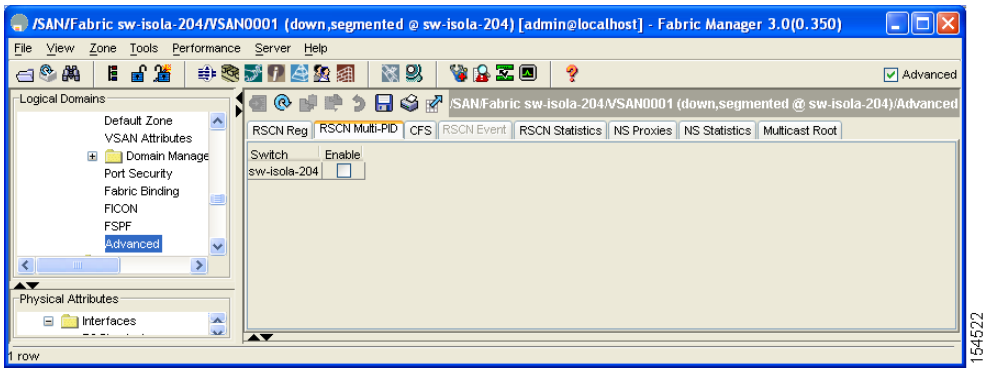

Step 3 Check the **Enable** check box.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to cancel any unsaved changes.

Clearing RSCN Statistics

You can clear the counters and later view the counters for a different set of events. For example, you can keep track of how many RSCNs or SW-RSCNs are generated on a particular event (like ONLINE or OFFLINE events). You can use these statistics to monitor responses for each event in the VSAN.

To clear the RSCN statistics for the specified VSAN, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

Configuring RSCN Timer Distribution Using CFS

Because the timeout value for each switch is configured manually, a misconfiguration occurs when different switches time out at different times. This means different N-ports in a network can receive RSCNs at different times. Cisco Fabric Services (CFS) alleviates this situation by automatically distributing configuration information to all switches in a fabric. This also reduces the number of SW-RSCNs.

RSCN supports two modes, distributed and nondistributed. In distributed mode, RSCN uses CFS to distribute configuration to all switches in the fabric. In nondistributed mode, only the configuration commands on the local switch are affected.

Note All configuration commands are not distributed. Only the **rscn event-tov** ***tov*** **vsan** ***vsan*** command is distributed.

The RSCN timer is registered with CFS during initialization and switchover. For high availability, if the RSCN timer distribution crashes and restarts or a switchover occurs, it resumes normal functionality from the state prior to the crash or switchover.

Note Before performing a downgrade, make sure that you revert the RCSN timer value in your network to the default value. Failure to do so will disable the links across your VSANs and other devices.

Compatibility across various Cisco MDS SAN-OS releases during an upgrade or downgrade is supported by **conf-check** provided by CFS. If you attempt to downgrade from Cisco MDS SAN-OS Release 3.0, you are prompted with a **conf-check** warning. You are required to disable RSCN timer distribution support before you downgrade.

By default, the RSCN timer distribution capability is disabled and is therefore compatible when upgrading from any Cisco MDS SAN-OS release earlier to 3.0.

Configuring the RSCN timer with CFS

To configure the RSCN timer with CFS using Fabric Manager, follow these steps:

Step 1 Expand a fabric, expand a VSAN and then select **Advanced** in the Logical Domains pane.

Step 2 Select the **RSCN Event** tab. (If it is greyed out, click the **CFS** tab first.)

You see the VSAN advanced configuration in the Information pane shown in Figure 29-6.

Figure 29-6 VSAN Advanced Configuration

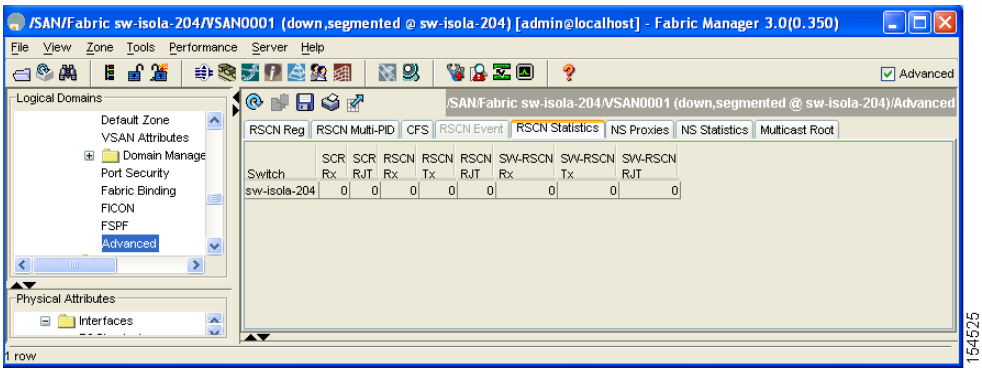

Step 3 Double-click the **TimeOut** value to change the value (in milliseconds) for the selected VSAN (see Figure 29-6).

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to cancel any unsaved changes.

Default Settings

Table 29-1 lists the default settings for RSCN.

Table 29-1 Default RSCN Settings

Parameters	Default
RSCN timer value	2000 milliseconds for Fibre Channel VSANs 1000 milliseconds for FICON VSANs
RSCN timer configuration distribution	Disabled

CHAPTER 30

Discovering SCSI Targets

This chapter describes the SCSI LUN discovery feature provided in switches in the Cisco MDS 9000 Family. It includes the following section:

About SCSI LUN Discovery

Small Computer System Interface (SCSI) targets include disks, tapes, and other storage devices. These targets do not register logical unit numbers (LUNs) with the name server.

The name server requires LUN information for the following reasons:

- To display LUN storage device information so an NMS can access this information.
- To report device capacity, serial number, and device ID information.
- To register the initiator and target features with the name server.

The SCSI LUN discovery feature uses the local domain controller Fibre Channel address. It uses the local domain controller as the source FC ID, and performs SCSI INQUIRY, REPORT LUNS, and READ CAPACITY commands on SCSI devices.

The SCSI LUN discovery feature is initiated on demand, through CLI or SNMP. This information is also synchronized with neighboring switches, if those switches belong to the Cisco MDS 9000 Family.

This section includes the following topics:

About Starting SCSI LUN Discovery

SCSI LUN discovery is done on demand.

Only Nx ports that are present in the name server database and that are registered as FC4 Type = SCSI_FCP are discovered.

Starting SCSI LUN Discovery

To begin SCSI LUN discovery using Device Manager, follow these steps:

Step 1 Choose **FC > Advanced > LUN**s.

You see the LUN Configuration dialog box.

Figure 30-1 LUN Configuration Dialog Box

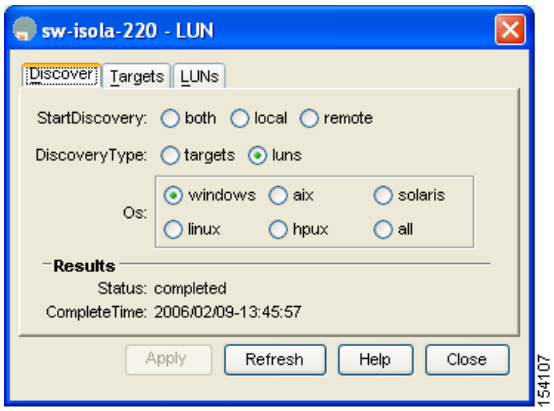

Step 2 Set StartDiscovery to **local**, **remote** or **both**.

Step 3 Choose the **DiscoveryType** and **OS**.

Step 4 Click **Apply** to begin discovery.

About Initiating Customized Discovery

Customized discovery consists of a list of VSAN and domain pairs that are selectively configured to initiate a discovery. The domain ID is a number from 0 to 255 in decimal or a number from 0x0 to 0xFF in hex.

Initiating Customized Discovery

To initiate a customized discovery using Device Manager, follow these steps:

Step 1 Click the VSAN drop-down menu and select the VSAN in which you want to initiate a customized discovery.

Step 2 Click **FC > Advanced > LUNs**.

You see the LUN Configuration dialog box.

Step 3 Set StartDiscovery to **local**, **remote** or **both**.

Step 4 Fill in the **DiscoveryType** and **OS** fields.

Step 5 Click **Apply** to begin discovery.

Displaying SCSI LUN Information

To display the results of the discovery using Device Manager, follow these steps:

Step 1 Click **FC > Advanced > LUN**s

You see the LUN Configuration dialog box.

Step 2 Click the **LUN** tab or the **Targets** tab.

CHAPTER 31

Configuring FICON

Fiber Connection (FICON) interface capabilities enhance the Cisco MDS 9000 Family by supporting both open systems and mainframe storage network environments. Inclusion of Control Unit Port (CUP) support further enhances the MDS offering by allowing in-band management of the switch from FICON processors.

The fabric binding feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations (see Chapter 42, “Configuring Fabric Binding”). The Registered Link Incident Report ((RLIR) application provides a method for a switch port to send an LIR to a registered Nx port.

This chapter includes the following sections:

About FICON

The Cisco MDS 9000 Family supports the Fibre Channel Protocol (FCP), FICON, iSCSI, and Fiber Channel over IP (FCIP) capabilities within a single, high availability platform. This solution simplifies purchasing, reduces deployment and management costs, and reduces the complex evolution to shared mainframe and open systems storage networks (see Figure 31-1).

Figure 31-1 Shared System Storage Network

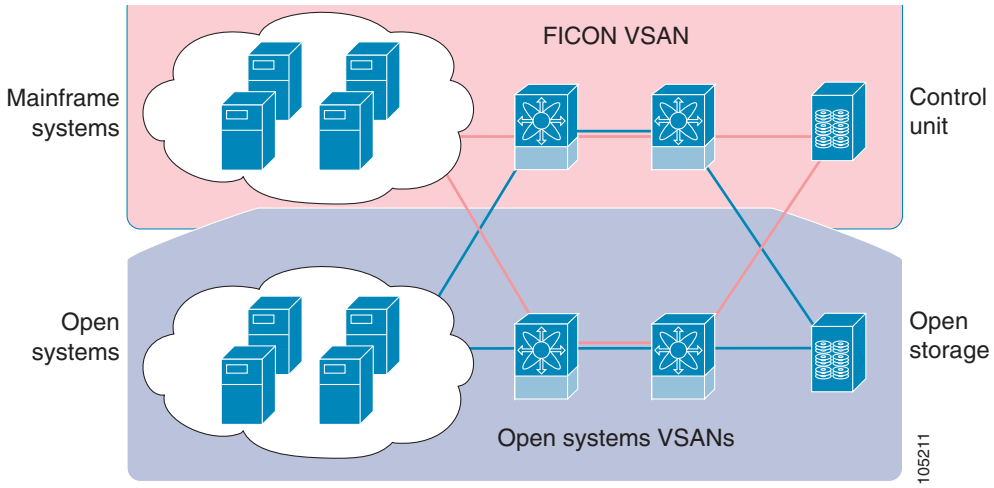

FCP and FICON are different FC4 protocols and their traffic is independent of each other. If required, devices using these protocols can be isolated using VSANs.

This section includes the following topics:

FICON Requirements

The FICON feature has the following requirements:

- You need the following switches to implement FICON features:
  - Any switch in the Cisco MDS 9500 Series.
  - Any switch in the Cisco MDS 9200 Series.

Note The FICON feature is not supported on Cisco MDS 9120 and 9140 switches or the 32-port Fibre Channel switching module.

- You need the MAINFRAME_PKG license to configure FICON parameters (see the "Obtaining and Installing Licenses" section on page 10-1).

MDS-Specific FICON Advantages

This section explains the additional FICON advantages in Cisco MDS switches.

Fabric Optimization with VSANs

Generally, separate physical fabrics have a high level of switch management and have a higher implementation cost. Further, the ports in each island may be over-provisioned depending on the fabric configuration.

By using the Cisco MDS-specific VSAN technology, you can introduce greater efficiency between these physical fabrics by lowering the cost of over-provisioning and reducing the number of switches to be managed.

VSANs also help you to move unused ports nondisruptively and provide a common redundant physical infrastructure (see Figure 31-2).

Figure 31-2 VSAN-Specific Fabric Optimization

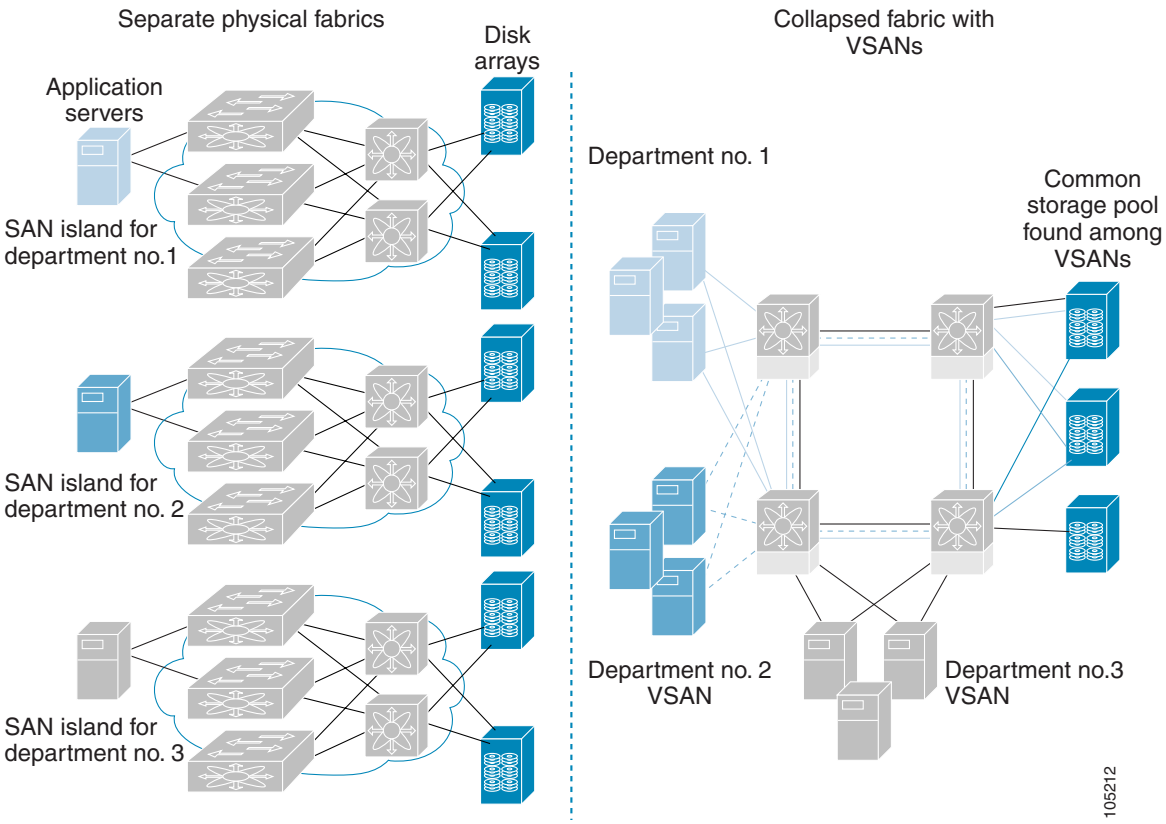

VSANs enable global SAN consolidation by allowing you to convert existing SAN islands into virtual SAN islands on a single physical network. It provides hardware-enforced security and separation between applications or departments to allow coexistence on a single network. It also allows virtual rewiring to consolidate your storage infrastructure. You can move assets between departments or applications without the expense and disruption of physical relocation of equipment.

Note While you can configure up to 256 VSANs in any Cisco MDS switch, you can enable FICON in eight of these VSANs.

FCIP Support

The multilayer architecture of the Cisco MDS 9000 Family enables a consistent feature set over a protocol-agnostic switch fabric. Cisco MDS 9500 Series and 9200 Series switches transparently integrate Fibre Channel, FICON, and Fibre Channel over IP (FCIP) in one system. The FICON over FCIP feature enables cost-effective access to remotely located mainframe resources. With the Cisco MDS 9000 Family platform, storage replication services such as IBM PPRC and XRC can be extended over metro to global distances using ubiquitous IP infrastructure and simplified business continuance strategies.

Caution When write-acceleration is enabled in an FCIP interface, a FICON VSAN will not be enabled in that interface. Likewise, if a FCIP interface is up in a FICON VSAN, write-acceleration cannot be enabled on that interface.

See Chapter 43, “Configuring FCIP.”

PortChannel Support

The Cisco MDS implementation of FICON provides support for efficient utilization and increased availability of Inter-Switch Links (ISLs) necessary to build stable large-scale SAN environments. PortChannels ensure an enhanced ISL availability and performance in Cisco MDS switches.

See the “Configuring PortChannels” section on page 21-1.

VSANs for FICON and FCP Mixing

Cisco MDS 9000 Family FICON-enabled switches simplify deployment of even the most complex mixed environments. Multiple logical FICON, Z-Series Linux/FCP, and Open-Systems FCP fabrics can be overlaid onto a single physical fabric by simply creating VSANs as required for each service. VSANs provide both hardware isolation and protocol-specific fabric services, eliminating the complexity and potential instability of zone-based mixed schemes.

By default, the FICON feature is disabled in all switches in the Cisco MDS 9000 Family. When the FICON feature is disabled, FC IDs can be allocated seamlessly. Intermixed environments are addressed by the Cisco SAN-OS software. The challenge of mixing FCP and FICON protocols are addressed by Cisco MDS switches when implementing VSANs.

Switches and directors in the Cisco MDS 9000 Family support FCP and FICON protocol mixing at the port level. If these protocols are mixed in the same switch, you can use VSANs to isolate FCP and FICON ports.

Tip When creating a mixed environment, place all FICON devices in one VSAN (other than the default VSAN) and segregate the FCP switch ports in a separate VSAN (other than the default VSAN). This isolation ensures proper communication for all connected devices.

Cisco MDS-Supported FICON Features

The Cisco MDS 9000 Family FICON features include:

- Flexibility and investment protection—The Cisco MDS 9000 Family shares common switching and service modules across the Cisco MDS 9500 Series and the 9200 Series.

  Refer to the *Cisco MDS 9500 Series Hardware Installation Guide* and the *Cisco MDS 9200 Series Hardware Installation Guide.*

- High-availability FICON-enabled director—The Cisco MDS 9500 Series combines nondisruptive software upgrades, stateful process restart and failover, and full redundancy of all major components for a new standard in director-class availability. It supports up to 224 autosensing, 2/1-Gbps, FICON or FCP ports in any combination in a single chassis and up to 768 Fibre Channel ports in a single rack. The 1.44 Tbps of internal system bandwidth ensures smooth integration of future 10-Gbps modules. See Chapter 15, "Configuring High Availability."

- Infrastructure protection—Common software releases provide infrastructure protection is available across all Cisco MDS 9000 platforms. See Chapter 13, "Software Images."

- VSAN technology—The Cisco MDS 9000 Family introduces VSAN technology for hardware-enforced, isolated environments within a single physical fabric for secure sharing of physical infrastructure and enhanced FICON mixed support. See Chapter 23, "Configuring and Managing VSANs."

- Port-level configurations—Each port has BB_credits, beacon mode, and port security for each port. See the "About Frame Encapsulation" section on page 18-7.

- Alias name configuration—Provides user-friendly alias names alias instead of the WWN for switches and attached node devices. See Chapter 26, "Configuring and Managing Zones."

- Comprehensive security framework—The Cisco MDS 9000 Family supports RADIUS authentication, Simple Network Management Protocol Version 3 (SNMPv3), role-based access control, Secure Shell Protocol (SSH), Secure File Transfer Protocol (SFTP), VSANs, hardware-enforced zoning, ACLs, fabric binding, Fibre Channel Security Protocol (FC-SP), LUN zoning, read-only zones, and VSAN-based access control. See Chapter 35, "Configuring RADIUS and TACACS+"and Chapter 40, "Configuring FC-SP and DHCHAP."

- Traffic encryption—IPSec is supported over FCIP. You can encrypt FICON and Fibre Channel traffic that is carried over FCIP. See Chapter 39, "Configuring IPsec Network Security."

- Local accounting log—View the local accounting log to locate FICON events. See the "MSCHAP Authentication" section on page 35-26.

- Unified storage management—Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for in-band management using the IBM S/A OS/390 I/O operations console. See the "CUP In-Band Management" section on page 31-35.

- Port address-based configurations—Configure port name, blocked or unblocked state, and the prohibit connectivity attributes. See the "FICON Ports" section on page 31-25.

- You can display the following information:

  - Individual Fibre Channel ports, such as the port name, port number, Fibre Channel address, operational state, type of port, and login data.
  - Nodes attached to ports.
  - Port performance and statistics.

  See the "Receiving FICON Alerts" section on page 31-37.

- Configuration files—Store and apply configuration files. See the "FICON Information Refresh Note" section on page 31-23.
- FICON and Open Systems Management Server features if installed. See the "VSANs for FICON and FCP Mixing" section on page 31-4.
- Enhanced cascading support—See the "CUP In-Band Management" section on page 31-35.
- Date and Time—Set the date and time on the switch. See the "About Host Control of the Time Stamp" section on page 31-22.
- Configure SNMP trap recipients and community names—See the "About SNMP Control of FICON Parameters" section on page 31-23.
- Call Home configurations—Configure the director name, location, description, and contact person. See Chapter 58, "Configuring Call Home."
- Configure preferred domain ID, FC ID persistence, and principle switch priority—See Chapter 22, "Configuring Domain Parameters."
- Sophisticated SPAN diagnostics—The Cisco MDS 9000 Family provides industry-first intelligent diagnostics, protocol, decoding, and network analysis tools as well as integrated Call Home capability for added reliability, faster problem resolution, and reduced service costs. See Chapter 56, "Monitoring Network Traffic Using SPAN."
- Configure R_A_TOV, E_D_TOV—See the "Common Interface Configuration" section on page 32-1.
- Director-level maintenance tasks—Perform maintenance tasks for the director including maintaining firmware levels, accessing the director logs, and collecting data to support failure analysis. See Chapter 64, "Monitoring System Processes and Logs."
- Port-level incident alerts—Display and clear port-level incident alerts.

FICON Cascading

The Cisco MDS SAN-OS software allows multiple switches in a FICON network. To configure multiple switches, you must enable and configure fabric binding in that switch (see the "Calculating FICON Flow Load Balance" section on page 31-39).

FICON VSAN Prerequisites

To ensure that a FICON VSAN is operationally up, be sure to verify the following requirements:

- Set the default zone to permit, if you are not using the zoning feature. See the "About the Default Zone" section on page 26-23.
- Enable in-order delivery on the VSAN. See Chapter 28, "Configuring Fibre Channel Routing Services and Protocols."
- Enable (and if required, configure) fabric binding on the VSAN. See the "Calculating FICON Flow Load Balance" section on page 31-39.
- Verify that conflicting persistent FC IDs do not exist in the switch. See Chapter 22, "Configuring Domain Parameters."

- Verify that the configured domain ID and requested domain ID match. See Chapter 22, "Configuring Domain Parameters."
- Add the CUP (area FE) to the zone, if you are using zoning. See the "CUP In-Band Management" section on page 31-35.

If any of these requirements are not met, the FICON feature cannot be enabled.

FICON Port Numbering

Default FICON port numbers are assigned by the Cisco MDS SAN-OS software based on the module and the slot in the chassis. The first port in a switch always starts with a zero (0) (see Figure 31-3).

Figure 31-3 Port Number in the Cisco MDS 9000 Family

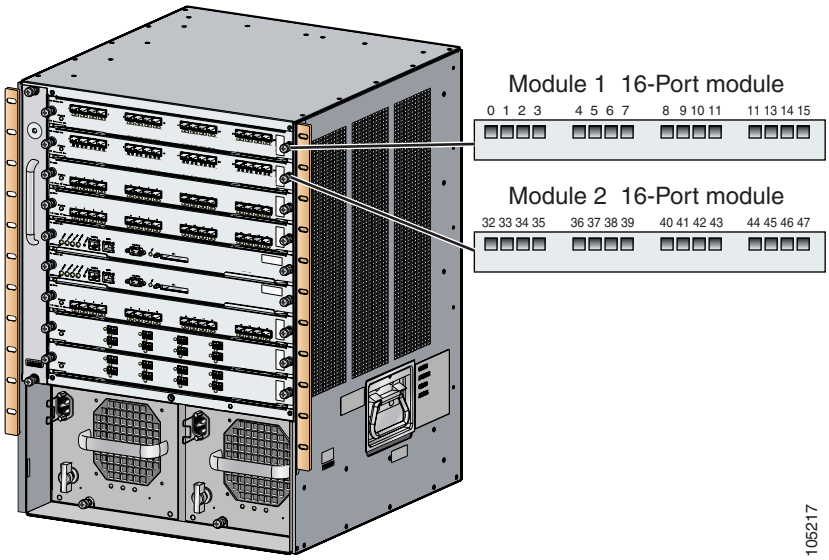

The default FICON port number is assigned based on the front panel location of the port and is specific to the slot in which the module resides. Thirty-two (32) port numbers are assigned to each slot on all Cisco MDS 9000 Family switches except for the Cisco MDS 9513 Director, which has 16 port numbers assigned for each slot. These default numbers are assigned regardless of the module's physical presence in the chassis, the port status (up or down), or the number of ports on the module (4, 12, 16, 24, or 48). If a module has fewer ports than the number of port numbers assigned to the slot, then the excess port numbers are unused. If a module has more ports than the number of port numbers assigned to the slot, the excess ports cannot be used for FICON traffic unless you manually assign the port numbers.

Note Follow the steps in "Assigning FICON Port Numbers to Slots" section on page 31-11 to make use of excess ports by manually assigning more port numbers to the slot. Before doing this, however, we recommend that you review the default port number assignments for Cisco MDS 9000 switches shown in Table 31-1 on page 31-9, and that you read the following sections to gain a complete understanding of FICON port numbering: "About the Reserved FICON Port Numbering Scheme" section on page 31-10, "FICON Port Numbering Guidelines" section on page 31-11, and "Assigning FICON Port Numbers to Slots" section on page 31-11.

Note Only Fibre Channel, PortChannel, and FCIP ports are mapped to FICON port numbers. Other types of interfaces do not have a corresponding port number.

This section includes the following topics:

Note You must enable FICON on the switch before reserving FICON port numbers (see the "About Enabling FICON" section on page 31-15).

FICON Port Number Assignment

The FICON port number is assigned based on the front panel location of the port and is specific to the slot in which the module resides. Even if the module is a 16-port module, 32 port numbers are assigned to that module—regardless of the module's physical presence in the chassis or the port status (up or down).

Note Only Fibre Channel, PortChannel, and FCIP ports are mapped to FICON port numbers. Other types of interfaces do not have a corresponding port number.

Table 31-1 lists the default port number assignment for the Cisco MDS 9000 Family of switches and directors.

Table 31-1 Default FICON Port Numbering in the Cisco MDS 9000 Family

Product	Slot Number	Implemented Port Allocation		Unimplemented Ports	Notes
		To Ports	To PortChannel/FCIP		
Cisco MDS 9200 Series	Slot 1	0 through 31	64 through 89	90 through 253 and port 255	Similar to a switching module.
	Slot 2	32 through 63			The first 4, 12, 16, or 24 port numbers in a 4-port, 12-port, 16-port, or 24-port module are used and the rest remain unused. Extra 16 ports on 48-port modules are not allocated numbers.
Cisco MDS 9506 Director	Slot 1	0 through 31	128 through 153	154 through 253 and port 255	
	Slot 2	32 through 63			
	Slot 3	64 through 95			
	Slot 4	96 through 127			
	Slot 5	None			Supervisor modules are not allocated port numbers.
	Slot 6	None			
Cisco MDS 9509 Director	Slot 1	0 through 31	224 through 249	250 through 253 and port 255	The first 4, 12, 16, or 24 port numbers in a 4-port, 12-port, 16-port, or 24-port module are used and the rest remain unused. Extra 16 ports on 48-port modules are not allocated port numbers.
	Slot 2	32 through 63			
	Slot 3	64 through 95			
	Slot 4	96 through 127			
	Slot 5	None			Supervisor modules are not allocated port numbers.
	Slot 6	None			
	Slot 7	128 through 159			The first 4, 12, 16, or 24 port numbers are used for a 4-port, 12-port,16-port, or 24-port module and the rest remain unused. Extra 16 ports on 48-port modules are not allocated port numbers.
	Slot 8	160 through 191			
	Slot 9	192 through 223			
Cisco MDS 9513 Director	Slot 1	0 through 15	191 through 226	227 through 253 and port 255	The first 4 or 12 port numbers are used for a 4-port or 12-port module and the rest remain unused.Extra ports on 24-port, 32-port, and 48-port modules are not allocated port numbers.
	Slot 2	16 through 31			
	Slot 3	32 through 63			
	Slot 4	64 through 79			
	Slot 5	80 through 95			
	Slot 6	96 through 111			
	Slot 7	None			Supervisor modules are not allocated port numbers.
	Slot 8	None			
	Slot 9	112 through 127			The first 4 or 12 port numbers are used for a 4-port or 12-port module and the rest remain unused. Extra ports on 24-port, 32-port, and 48-port modules are not allocated port numbers.
	Slot 10	128 through 143			
	Slot 11	144 through 159			
	Slot 12	160 through 175			
	Slot 13	176 through 191			

Port Addresses

By default, port numbers are the same as port addresses. You can swap the port addresses (see the "Port Swapping" section on page 31-33).

Implemented and Unimplemented Port Addresses

An implemented port refers to any port address that is available in the chassis (see Table 31-1). An unimplemented port refers to any port address that is not available in the chassis (see Table 31-1).

About the Reserved FICON Port Numbering Scheme

A range of 255 port numbers are available for you to assign to all the ports on a switch. Table 31-1 shows that you can have more than 255 physical ports on a switch and the excess ports do not have port numbers in the default numbering scheme. When you have more than 255 physical ports on your switch, you can assign unimplemented port numbers to the ports, or assign duplicate port numbers if they are not used in the same FICON VSAN. For example, you can configure port number 1 on interface fc1/1 in FICON VSAN 10 and fc10/1 in FICON VSAN 20.

Note A VSAN can have a maximum of 250 port numbers.

Note FICON port numbers are not changed for ports that are active. You must first disable the interfaces using the **shutdown** command.

Note You can configure port numbers even when no module is installed in the slot.

Installed and Uninstalled Ports

An installed port refers to a port for which all required hardware is present. A specified port number in a VSAN can be implemented, and yet not installed, if any of the following conditions apply:

- The module is not present—For example, if module 1 is not physically present in slot 1 in a Cisco MDS 9509 Director, ports 0 to 31 are considered uninstalled.
- The small form-factor pluggable (SFP) port is not present—For example, if a 16-port module is inserted in slot 2 in a Cisco MDS 9509 Director, ports 48 to 63 are considered uninstalled.
- The port is not in a FICON-enabled VSAN—For example, if port 4 (of a 16-port module in slot 1) is configured in FICON-enabled VSAN 2, then only port 4 is installed and ports 0 to 3 and 5 to 15 are uninstalled—even if they are implemented in VSAN 2.

  Another scenario is if VSANs 1 through 5 are FICON-enabled, and trunking-enabled interface fc1/1 has VSANs 3 through 10, then port address 0 is uninstalled in VSAN 1 and 2.
- The port is part of a PortChannel—For example, if interface fc 1/1 is part of PortChanne1 5, port address 0 is uninstalled in all FICON VSANs. See Table 31-1.

FICON Port Numbering Guidelines

The following guidelines apply to FICON port numbers:

- Supervisor modules do not have port number assignments.
- Port numbers are VSAN independent and do not change based on VSANs or TE ports.
- Each PortChannel must be explicitly associated with a FICON port number.
- When the port number for a physical PortChannel becomes uninstalled, the relevant PortChannel configuration is applied to the physical port.
- Each FCIP tunnel must be explicitly associated with a FICON port number. If the port numbers are not assigned for PortChannels or for FCIP tunnels, the associated ports will not come up.

  See the "About Port Numbers for FCIP and PortChannel Interfaces" section on page 31-12.

Assigning FICON Port Numbers to Slots

To assign FICON port numbers to slots using Device Manager, follow these steps:

Step 1 Click **FICON** and then select **Port Numbers**. You see the FICON port numbers (see Figure 31-4).

Figure 31-4 FICON Port Numbers

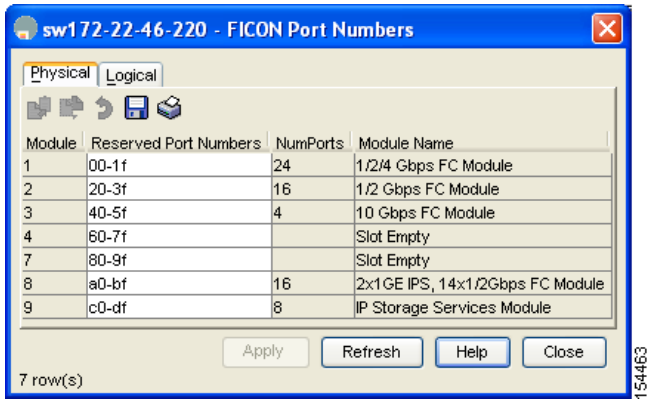

Step 2 Enter the chassis slot port numbers in the Reserved Port Numbers field.

Step 3 Click **Apply**.

Step 4 Click **Close**.

About Port Numbers for FCIP and PortChannel Interfaces

FCIP and PortChannels cannot be used in a FICON-enabled VSAN unless they are explicitly bound to a port number.

See the "FICON Ports" section on page 31-25 and the "Reserving FICON Port Numbers for FCIP and PortChannel Interfaces" section on page 31-13.

You can use the default port numbers if they are available (see Table 31-1 on page 31-9) or if you reserve port numbers from the pool of port numbers that are not reserved for Fibre Channel interfaces (see the "FICON Port Numbering" section on page 31-7).

Reserving FICON Port Numbers for FCIP and PortChannel Interfaces

You must reserve port numbers for logical interfaces, such as FCIP and PortChannels, if you plan to use them.

To reserve FICON port numbers for FCIP and PortChannel interfaces using Device Manager, follow these steps:

Step 1 Click **FICON > Port Numbers.** You see the FICON port numbers.

Step 2 Click the **Logical** tab to see the reserved port numbers for the slot (see Figure 31-5).

Figure 31-5 Reserved Port Numbers for the Selected Slot

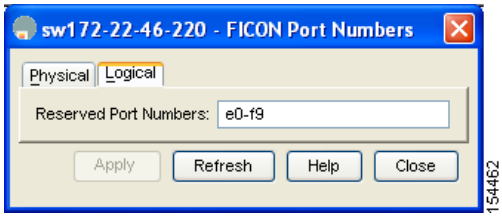

Step 3 Enter the chassis slot port numbers (see Figure 31-5). These are the reserved port numbers for one chassis slot. There can be up to 64 port numbers reserved for each slot in the chassis.

Step 4 Click **Apply**.

Step 5 Click **Close**.

FC ID Allocation

FICON requires a predictable and static FC ID allocation scheme. When FICON is enabled, the FC ID allocated to a device is based on the port address of the port to which it is attached. The port address forms the middle byte of the fabric address. Additionally, the last byte of the fabric address should be the same for all devices in the fabric. By default, the last byte value is 0 and can be configured (see the "About FC ID Last Byte" section on page 31-20).

Note You cannot configure persistent FC IDs in FICON-enabled VSANs.

Cisco MDS switches have a dynamic FC ID allocation scheme. When FICON is enabled or disabled on a VSAN, all the ports are flapped to switch from the dynamic to static FC IDs and vice versa (see Figure 31-6).

Figure 31-6 Static FC ID Allocation for FICON

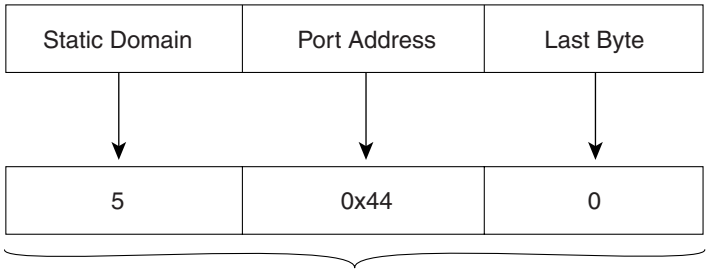

Static FC ID allocation for interface fc3/5 includes the static domain ID (5), the port address (0x44), and the last byte value (0).

FICON Configuration

By default FICON is disabled in all switches in the Cisco MDS 9000 Family. You can enable FICON on a per VSAN basis by using the Device Manager.

This section includes the following topics:

About Enabling FICON

When you enable the FICON feature in Cisco MDS switches, the following apply:

- You cannot disable in-order delivery for the FICON-enabled VSAN.
- You cannot disable fabric binding or static domain ID configurations for the FICON-enabled VSAN.
- The load balancing scheme is changed to Source ID (SID)—Destination ID (DID). You cannot change it back to SID—DID—OXID.
- The IPL configuration file is automatically created.

See the "FICON Information Refresh Note" section on page 31-23.

See the "About FICON Configuration Files" section on page 31-31.

Enabling FICON

By default FICON is disabled in all switches in the Cisco MDS 9000 Family. You can enable FICON on the switch either explicitly or implicitly by enabling FICON on a VSAN. However, disabling FICON on all VSANs does not disable FICON on the switch. You must explicitly disable FICON.

Tip This section describes the procedure to manually enable FICON on a VSAN. If you have already enabled FICON on the required VSAN using the automated setup (recommended), skip to the "FICON Information Refresh Note" section on page 31-23.

Note Using Device Manager, FICON auto-save can be invoked by multiple users logged on to the same FICON-enabled switch. Device Manager performs a periodic auto-save on any FICON-enabled switch causing increments in the FICON key counter. These increments highlight a change that has actually not occurred. To avoid this we recommend that only one instance of Device Manager monitor a FICON-enabled switch.

To create a FICON-enabled VSAN using Fabric Manager, follow these steps:

Step 1 Click the **Create VSAN** icon (see Figure 31-7).

Figure 31-7 Create VSAN Icon

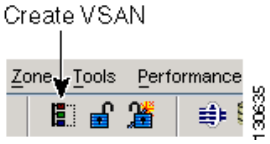

You see the Create VSAN dialog box (see Figure 31-8).

Figure 31-8 Create VSAN Dialog Box

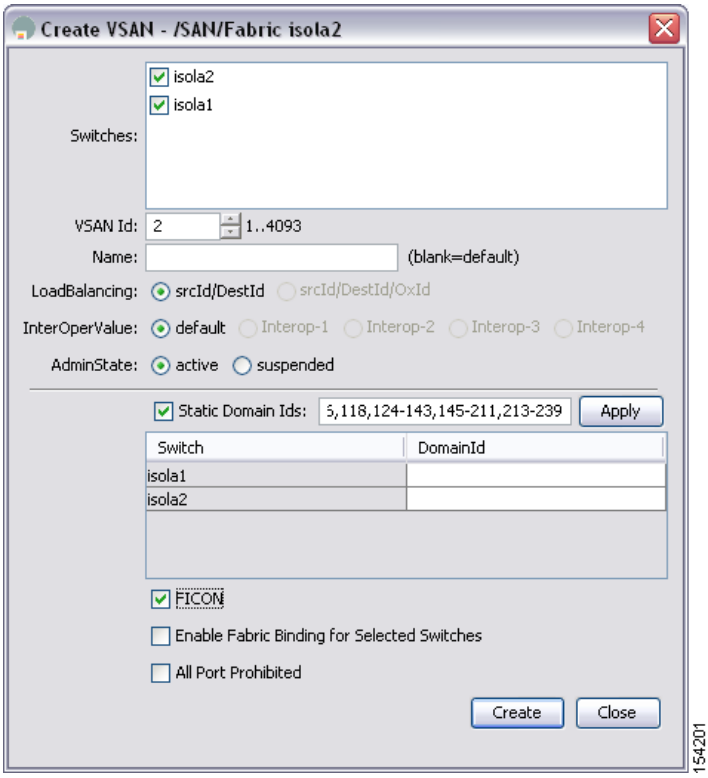

Step 2 Select the switches you want to be in the VSAN.

Step 3 Enter a VSAN ID.

Step 4 Enter the name of the VSAN, if desired.

Step 5 Select the type of load balancing, the interop value, and the administrative state for this VSAN.

Step 6 Check the **FICON** check box.

Note You cannot enable interop modes on FICON-enabled VSANs.

Step 7 Check the option, if appropriate, to enable fabric binding for the selected switches.

Step 8 Check the All Ports Prohibited option if all ports in this VSAN are prohibited.

Step 9 Click **Create** to create the VSAN, or click **Close** to close the dialog without creating the VSAN.

Step 10 Open Device Manager for each switch in the FICON VSAN by clicking **Tools** and selecting **Device Manager**.

Step 11 Click **VSANs** from the FC menu.

You see the VSAN dialog box (see Figure 31-9).

Figure 31-9 VSAN Dialog Box in Device Manager

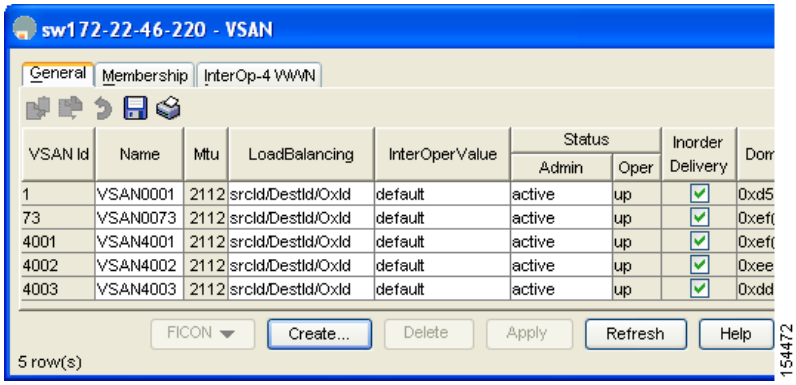

Step 12 Enter the VSAN membership information.

Step 13 Click the VSAN you want to become a FICON VSAN and select **Add** from the FICON drop-down menu.

Step 14 Click **Apply** to save these changes or click **Close** to exit the dialog box without saving changes.

Manually Enabling FICON on a VSAN

To manually enable FICON on a VSAN using Fabric Manager, follow these steps:

Step 1 Click Fabric > VSAN > FICON.

You see the FICON VSAN configuration information in the Information pane.

Step 2 Select the switch in the VSAN on which you want to enable FICON.

Step 3 Click enable from the Command drop-down menu.

Step 4 Click **Apply Changes**.

Deleting FICON VSANs

To delete a FICON VSAN using Fabric Manager, follow these steps:

Step 1 Select **All VSANS**.

You see the VSAN table in the Information pane (see Figure 31-10).

Figure 31-10 All VSANs Table

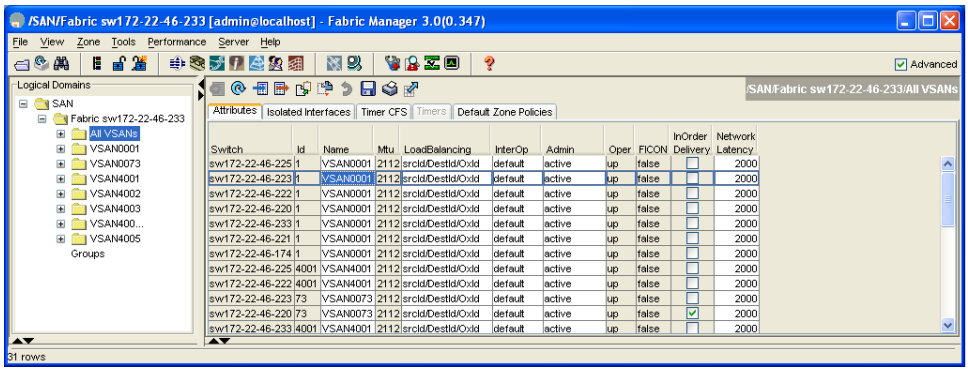

Step 2 Click anywhere in the row of the VSAN that you want to delete.

Step 3 Click **Delete Row** to delete the VSAN.

Note Deleting the VSAN will also delete the associated FICON configuration file, and the file cannot be recovered.

Suspending a FICON VSAN

To suspend a FICON VSAN using Fabric Manager, follow these steps:

Step 1 Click **All VSANs**.

You see all the VSANs listed in the Information pane.

Step 2 Select the VSAN that you want to suspend.

Step 3 Set the Admin drop-down menu for a VSAN to **suspended**.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

Note This command can be issued by the host if the host is allowed to do so (see the "Allowing the Host to Move the Switch Offline" section on page 31-20).

About the code-page Option

FICON strings are coded in Extended Binary-Coded Decimal Interchange Code (EBCDIC) format. Refer to your mainframe documentation for details on the code page options.

Cisco MDS switches support **international-5**, **france**, **brazil**, **germany**, **italy**, **japan**, **spain-latinamerica**, **uk**, and **us-canada** (default) EBCDIC format options.

Tip This is an optional configuration. If you are not sure of the EBCDIC format to be used, we recommend retaining the **us-canada** (default) option.

Configuring the code-page Option

To modify the code-page option using Device Manager, follow these steps:

Step 1 Select **VSANs** from the FICON menu.

You see the FICON VSAN configuration dialog box (see Figure 31-11). The VSANs tab is the default tab.

Figure 31-11 FICON VSANs Tab in Device Manager

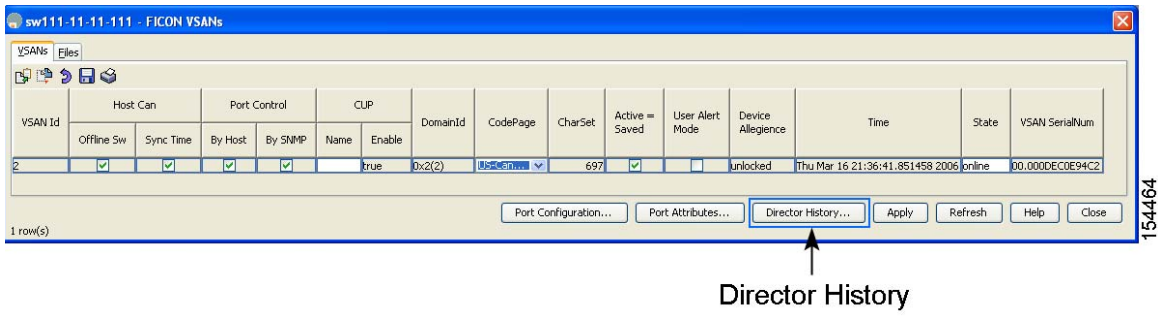

Step 2 Choose an option from the CodePage drop-down menu for the FICON VSAN you want to configure (US-Canada is configured in Figure 31-11).

Step 3 Click **Apply** to save the changes or click **Close** to exit the dialog box without saving changes.

About FC ID Last Byte

Caution If the FICON feature is configured in cascaded mode, the Cisco MDS switches use ISLs to connect to other switches.

FICON requires the last byte of the fabric address to be the same for all allocated FC IDs. By default, this value is set to 0. You can only change the FC ID last byte when the FICON switch is in the offline state.

See the "CUP In-Band Management" section on page 31-35.

Assigning the FC ID Last Byte

To assign the last byte for the FC ID using Fabric Manager, follow these steps:

Step 1 Click Fabric > VSAN. Click **Domain Manager**.

Step 2 Click the **Persistent FCIDs** tab.

You see the Persistent FcIds tab (see Figure 31-12).

Figure 31-12 Persistent Fclds Tab

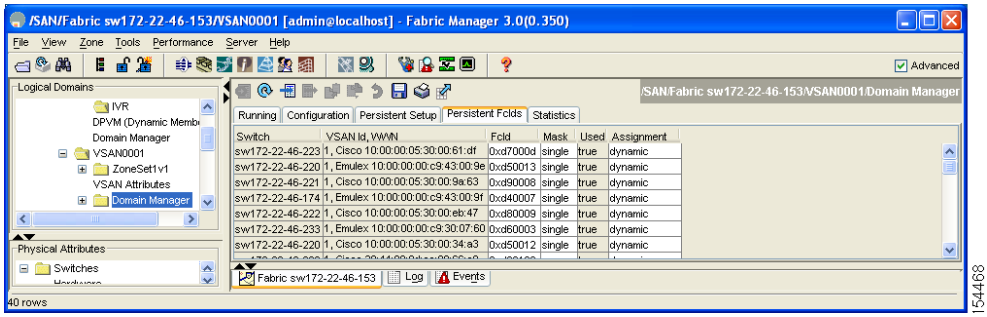

Step 3 Select **single** in the Mask column and then assign the entire FC ID at once. The single option allows you to enter the FC ID in the ###### format.

Step 4 Click **Apply Changes**.

Allowing the Host to Move the Switch Offline

By default, hosts are allowed to move the switch to an offline state.

To allow the host (mainframe) to move the switch to an offline state using Fabric Manager, follow these steps:

Step 1 Click Fabric > VSAN. Select **FICON**.

You see a list of switches under the Control tab in the Information pane.

Step 2 Click the **VSANs** tab.

You see the FICON VSAN configuration information in the Information pane (see Figure 31-13).

Figure 31-13 FICON VSANs in Fabric Manager

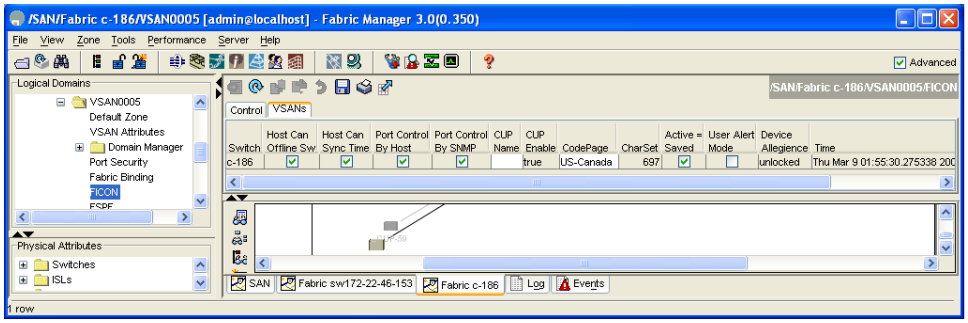

Step 3 Check the **Host Can Offline Sw** checkbox to allow the mainframe to move a switch to the offline state (see Figure 31-13).

Step 4 Check the **Host Can Sync Time** checkbox to allow the mainframe to set the system time on the switch (see Figure 31-13).

Step 5 Click **Apply Changes** to save the changes or click **Undo Changes** to discard any unsaved changes.

Allowing the Host to Change FICON Port Parameters

By default, mainframe users are not allowed to configure FICON parameters on Cisco MDS switches—they can only query the switch.

To allow the host (mainframe) to configure FICON parameters on the Cisco MDS switch using Fabric Manager, follow these steps:

Step 1 Click Fabric > VSAN. Select **FICON**.

You see a list of switches under the **Control** tab in the Information pane.

Step 2 Click the **VSANs** tab.

You see the FICON VSAN configuration information in the Information pane (see Figure 31-14).

Figure 31-14 FICON VSANs in Fabric Manager

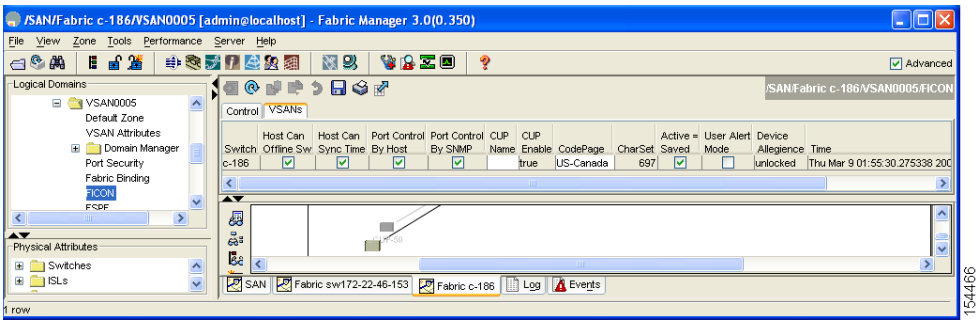

Step 3 Check the **Port Control By Host** checkbox to allow the mainframe to control a switch.

Step 4 Click **Apply Changes** to save the changes or click **Undo Changes** to discard any unsaved changes.

About Host Control of the Time Stamp

By default, the clock in each VSAN is the same as the switch hardware clock. Each VSAN in a Cisco MDS 9000 Family switch represents a virtual director. The clock and time present in each virtual director can be different.To maintain separate clocks for each VSAN, the Cisco SAN-OS software maintains the difference of the VSAN-specific clock and the hardware-based director clock. When a host (mainframe) sets the time, the Cisco SAN-OS software updates this difference between the clocks. When a host reads the clock, it computes the difference between the VSAN-clock and the current director hardware clock and presents a value to the mainframe.

Allowing the Host to Control the Time Stamp

To configure host (mainframe) control for the VSAN time stamp using Fabric Manager, follow these steps:

Step 1 Click Fabric > VSAN. Select **FICON**.

You see a list of switches under the Control tab in the Information pane.

Step 2 Click the **VSANs** tab.

You see the FICON VSAN configuration information in the Information pane (see Figure 31-15).

Figure 31-15 FICON VSANs in Fabric Manager

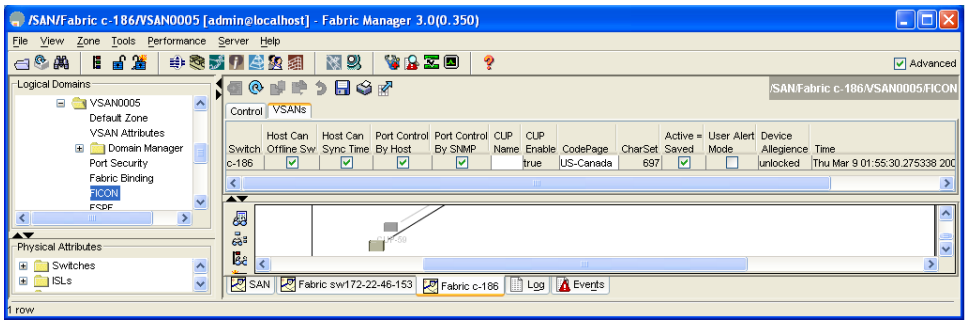

Step 3 Check the **Host Can Sync Time** checkbox to allow the mainframe to set the system time on the switch.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

About SNMP Control of FICON Parameters

By default, SNMP users can configure FICON parameters through the Cisco MDS 9000 Family Fabric Manager.

Note If you disable SNMP in the Cisco MDS switch, you cannot configure FICON parameters using the Fabric Manager.

Configuring SNMP Control of FICON Parameters

To configure SNMP control of FICON parameters using Fabric Manager, follow these steps:

Step 1 Click Fabric > VSAN. Select **FICON**.

You see a list of switches under the Control tab in the Information pane.

Step 2 Click the **VSANs** tab.

You see the FICON VSAN configuration information in the Information pane.

Step 3 Check the **Port Control By SNMP** checkbox to allow SNMP users to configure FICON on the switch.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

FICON Information Refresh Note

When viewing FICON information through the Device Manager dialog boxes, you must manually refresh the display by clicking the **Refresh** button to see the latest updates. This is true whether you configure FICON through the CLI or through the Device Manager.

There is no automatic refresh of FICON information. This information would be refreshed so often that it would affect performance.

About FICON Device Allegiance

FICON requires that serialization of access among multiple mainframes, CLI, and SNMP sessions be maintained on Cisco MDS 9000 Family switches by controlling device allegiance for the currently executing session. Any other session is denied permission to perform configuration changes unless the required allegiance is available.

Caution This task discards the currently executing session.

About Automatically Saving the Running Configuration

Cisco MDS SAN-OS provides an option to automatically save any configuration changes to the startup configuration. This ensures that the new configuration is present after a switch reboot. The Active=Saved option can be enable on any FICON VSAN.

Table 31-2 displays the results of the Active = Saved option and the implicit copy from the running configuration to the startup configuration (**copy running start**) in various scenarios.

If the Active=Saved option is enabled in any FICON-enabled VSAN in the fabric, then the following apply (see Number 1 and 2 in Table 31-2):

- All configuration changes (FICON-specific or not) are automatically saved to persistent storage (implicit **copy running start**) and stored in the startup configuration.
- FICON-specific configuration changes are immediately saved to the IPL file (see the "FICON Information Refresh Note" section on page 31-23).

If the Active=Saved option is not enabled in any FICON-enabled VSAN in the fabric, then FICON-specific configuration changes are not saved in the IPL file and an implicit **copy running startup** is not issued—you must explicitly save the running configuration to the startup configuration (see number 3 in Table 31-2).

Table 31-2 Saving the Active FICON and Switch Configuration

Number	FICON-enabled VSAN?	Active = Saved Enabled?	Implicit[1] `copy running start` Issued?	Notes
1	Yes	Yes (in all FICON VSANs)	Implicit	FICON changes written to the IPL file. Non-FICON changes saved to startup configuration and persistent storage. **Note** Interop modes cannot be enbled on FICON-enabled VSANs.
2		Yes (even in one FICON VSAN)	Implicit	FICON changes written to IPL file for only the VSAN that has Active=Saved enabled. Non-FICON changes saved to startup configuration and persistent storage.
3		Not in any FICON VSAN	Not implicit	FICON changes are not written to the IPL file. Non-FICON changes are saved to startup configuration only if you explicitly save the running configuration to the startup configuration.
4	No	Not applicable		

1. When the Cisco SAN-OS software implicitly saves the running configuration to the startup configuration in the Cisco MDS switch, only a binary configuration is generated—an ASCII configuration is not generated. If you wish to generate an additional ASCII configuration at this stage, you must explicitly copy the running configuration to the startup configuration.

Automatically Saving the Running Configuration

To automatically save the running configuration, follow these steps:

Step 1 Click Fabric > VSAN. Select **FICON**.

You see a list of switches under the Control tab in the Information pane.

Step 2 Click the **VSANs** tab.

You see the FICON VSAN configuration information in the Information pane.

Step 3 Check the **Active=Saved** check box to automatically save the running configuration to the startup configuration whenever there is a FICON configuration change.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

FICON Ports

You can perform FICON configurations on a per-port address basis in the Cisco MDS 9000 Family of switches.

Even if a port is uninstalled, the port address-based configuration is accepted by the Cisco MDS switch. This configuration is applied to the port when the port becomes installed.

This section includes the following topics:

About Port Blocking

If you block a port, the port is retained in the operationally down state. If you unblock a port, a port initialization is attempted. When a port is blocked, data and control traffic are not allowed on that port.

Physical Fibre Channel port blocks will continue to transmit an off-line state (OLS) primitive sequence on a blocked port.

Caution You cannot block or prohibit the CUP port (0XFE).

If a port is shut down, unblocking that port does not initialize the port.

Note The **shutdown/no shutdown** port state is independent of the **block/no block** port state.

Configuring Port Blocking

To block or unblock port addresses in a VSAN using Device Manager, follow these steps:

Step 1 Click ***Fabricxxx*** > ***VSANxxx***. Select FICON.

You see the FICON configuration table in the Information pane (see Figure 31-16).

Step 2 Click the Port Configuration tab.

Figure 31-16 FICON Port Configuration Dialog Box

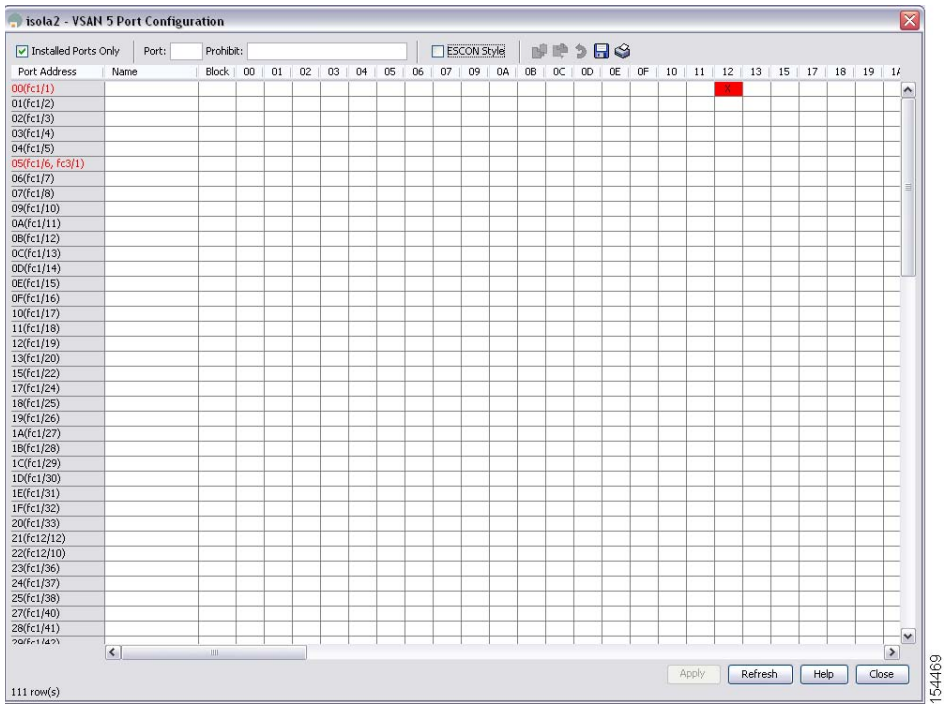

Step 3 Check the **Blocked** checkbox for the port that you want to block.

Step 4 Click **Apply** to save the changes or click **Close** to exit the dialog box without saving changes.

Viewing ESCON Style Ports

To view the available and prohibited ESCON style ports using Device Manager, follow these steps:

Step 1 Select **VSANs** from the FICON menu.

You see the FICON VSAN configuration dialog box. The **VSANs** tab is the default.

Step 2 Select a VSAN ID and click **Port Configuration**.

You see the FICON Port Configuration dialog box.

Step 3 Check the ESCON Style check box.

You can see the available and prohibited ESCON style ports. In Figure 31-17, A stands for available and P stands for prohibited.

When the port address is highlighted red, it represents the E/TE port or multiple interfaces.

Figure 31-17 ESCON Style

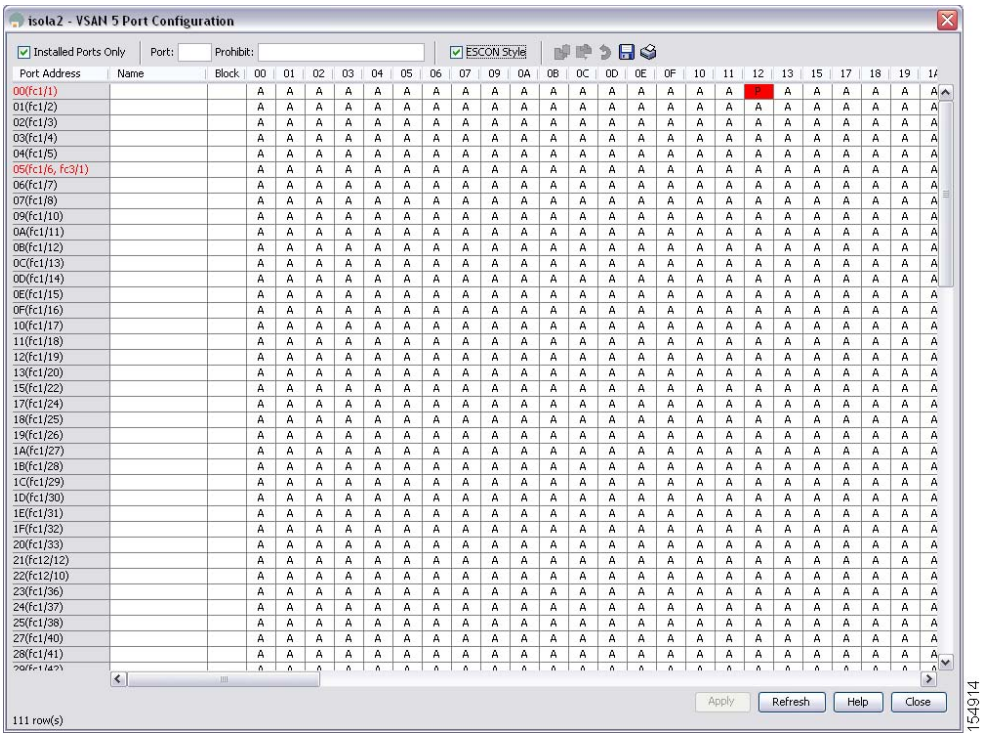

Step 4 Click **Apply** to save the changes or click **Close** to exit the dialog box without saving changes.

About Port Prohibiting

To prevent implemented ports from talking to each other, configure prohibits between two or more ports. If you prohibit ports, the specified ports are prevented from communicating with each other.

Tip You cannot prohibit a PortChannel or FCIP interface.

Unimplemented ports are always prohibited. In addition, prohibit configurations are always symmetrically applied—if you prohibit port 0 from talking to port 15, port 15 is automatically prohibited from talking to port 0.

Note If an interface is already configured in E or TE mode and you try to prohibit that port, your prohibit configuration is rejected. Similarly, if a port is not up and you prohibit that port, the port is not allowed to come up in E mode or in TE mode.

Configuring Port Prohibiting

To prohibit port addresses in a VSAN using Device Manager, follow these steps:

Step 1 Select **VSANs** from the FICON menu.

You see the FICON VSAN configuration dialog box. The **VSANs** tab is the default.

Step 2 Select a VASAN ID and click **Port Configuration**.

You see the FICON Port Configuration dialog box.

Step 3 Set the port prohibit configuration for the selected FICON VSANs.

Step 4 Click **Apply** to save these changes or click **Close** to exit the dialog box without saving changes.

Assigning a Port Address Name

Note To view the latest FICON information, you must click the **Refresh** button. See the "FICON Information Refresh Note" section on page 31-23.

To assign a port address name in Device Manager, follow these steps:

Step 1 Select **VSANs** from the FICON menu.

You see the FICON VSAN configuration dialog box. The **VSANs** tab is the default.

Step 2 Select a VSAN ID and click **Port Configuration**.

You see the FICON Port Configuration dialog box.

Step 3 Enter the Port Configuration information.

Step 4 Click **Apply** to save the configuration information or click **Cancel** to exit the dialog without saving.

About RLIR

The Registered Link Incident Report (RLIR) application provides a method for a switch port to send an Link Incident Record (LIR) to a registered Nxport. RLIR is a highly available application.

If an LIR is detected in FICON-enabled switches in the Cisco MDS 9000 Family from a RLIR Extended Link Service (ELS), it sends that record to the members in its Established Registration List (ERL).

In case of a multi-switch topology, a Distribute Registered Link Incident Record (DRLIR) Inter-Link Service (ILS) is sent to all reachable remote domains along with the RLIR ELS. On receiving the DRLIR ILS, the switch extracts the RLIR ELS and sends it to the members of the ERL.

The Nx-ports interested in receiving the RLIR ELS send Link Incident Record Registration (LIRR) ELS request to the management server on the switch. The RLIRs are processed on a per-VSAN basis.

The RLIR data is written to persistent storage when you copy the running configuration to the startup configuration.

Displaying RLIR Information

To view RLIR information using Device Manager, follow these steps:

Step 1 Choose **FICON > RLIR ERL**.

You see the Show RLIR ERL dialog box.

Step 2 Click **Close** to close the dialog box.

FICON Configuration Files

You can save up to 16 FICON configuration files on each FICON-enabled VSAN (in persistent storage). The file format is proprietary to IBM. These files can be read and written by IBM hosts using the in-band CUP protocol. Additionally, you can use the Cisco MDS CLI or Fabric Manager applications to operate these FICON configuration files.

Note Multiple FICON configuration files with the same name can exist in the same switch, provided they reside in different VSANs. For example, you can create a configuration file named XYZ in both VSAN 1 and VSAN 3.

When you enable the FICON feature in a VSAN, the switches always use the startup FICON configuration file, called IPL. This file is created with a default configuration as soon as FICON is enabled in a VSAN.

Caution When FICON is disabled on a VSAN, all the FICON configuration files are irretrievably lost.

FICON configuration files contain the following configuration for each implemented port address:

- Block
- Prohibit mask
- Port address name

Note Normal configuration files used by Cisco MDS switches include FICON-enabled attributes for a VSAN, port number mapping for PortChannels and FCIP interfaces, port number to port address mapping, port and trunk allowed VSAN configuration for ports, in-order guarantee, static domain ID configuration, and fabric binding configuration.

See the Chapter 11, “Initial Configuration,” for details on the normal configuration files used by Cisco MDS switches.

This section includes the following topics:

About FICON Configuration Files

Only one user can access the configuration file at any given time:

- If this file is being accessed by user 1, user 2 cannot access this file.
- If user 2 does attempt to access this file, an error is issued to user 2.
- If user 1 is inactive for more than 15 seconds, the file is automatically closed and available for use by any other permitted user.

FICON configuration files can be accessed by any host, SNMP, or CLI user who is permitted to access the switch. The locking mechanism in the Cisco SAN-OS software restricts access to one user at a time per file. This lock applies to newly created files and previously saved files. Before accessing any file, you must lock the file and obtain the file key. A new file key is used by the locking mechanism for each lock request. The key is discarded when the lock timeout of 15 seconds expires. The lock timeout value cannot be changed.

Applying the Saved Configuration Files to the Running Configuration

To apply the saved configuration files to the running configuration using Device Manager, follow these steps:

Step 1 Select **VSANs** from the FICON menu.

You see the FICON VSANs configuration dialog box. The **VSANs** tab is default.

Step 2 Click the **Files** tab.

You see the FICON Files dialog box (see Figure 31-18).

Figure 31-18 FICON VSANs Dialog Box

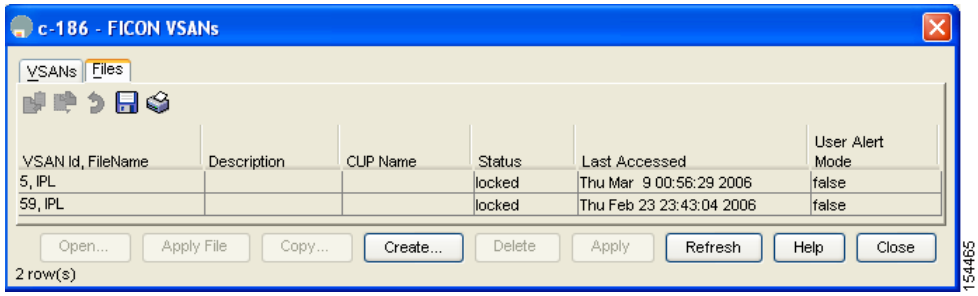

Step 3 Highlight the file you want to apply and click **Apply File** to apply the configuration to the running configuration.

About Editing FICON Configuration Files

Note To view the latest FICON information, you must click the **Refresh** button. See the "FICON Information Refresh Note" section on page 31-23.

The configuration file submode allows you to create and edit FICON configuration files. If a specified file does not exist, it is created. Up to 16 files can be saved. Each file name is restricted to eight alphanumeric characters.

Editing FICON Configuration Files

To edit the contents of a specified FICON configuration file using Device Manager, follow these steps:

Step 1 Select **VSANs** from the FICON menu.

You see the FICON VSAN configuration dialog box. The **VSANs** tab is default.

Step 2 Click the **Files** tab.

You see the FICON VSANs dialog box (see Figure 31-19).

Figure 31-19 FICON VSANs Dialog Box in Device Manager

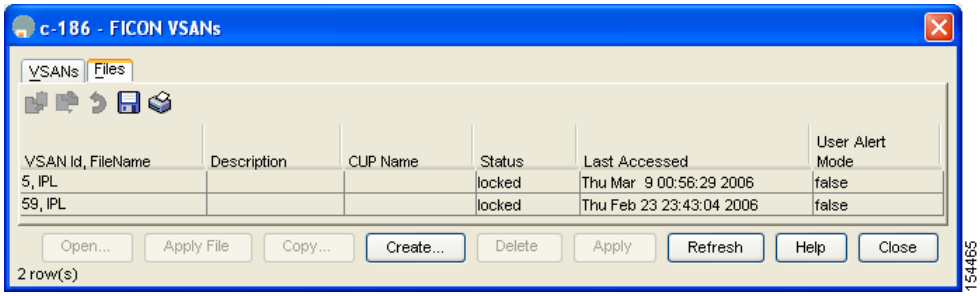

Step 3 Select a VSAN ID and then click **Open** to edit the FICON configuration file.

Step 4 Select a VSAN ID and then click **Delete** to delete the FICON configuration file.

Step 5 Click **Apply** to apply the changed FICON configuration file.

To open and view configuration files in Fabric Manager, follow these steps:

Step 1 Click Fabricxxx > VSANxxx. Select FICON.

You see the FICON configuration table in the Information pane.

Step 2 Click the Files tab.

Step 3 Select the file you want to open.

Step 4 Click Open.

Copying FICON Configuration Files

To copy an existing FICON configuration file using Device Manager, follow these steps:

Step 1 Select **VSANs** from the FICON menu.

You see the FICON VSAN configuration dialog box. The **VSANs** tab is default.

Step 2 Click the **Files** tab.

You see the FICON Files dialog box.

Step 3 Click **Create** to create a FICON configuration file.

You see the Create FICON Configuration File dialog box in Figure 31-20.

Figure 31-20 Create FICON Configuration File Dialog Box in Device Manager

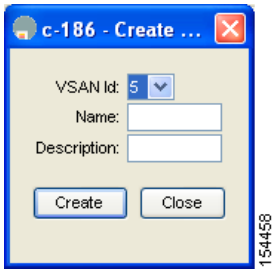

a. Select a VSAN ID for the FICON VSAN you want to configure.

b. Enter the file name and the description.

c. Click **Create** to create the file, or click **Close** to close the dialog without creating the file.

Step 4 Click **Copy** to copy the file to a new file.

Step 5 Click **Apply** to apply the FICON configuration file.

You can see the list of existing configuration files by clicking **FICON > VSANs** and selecting the **Files** tab.

To view the

Port Swapping

The FICON port swapping feature is only provided for maintenance purposes.

The FICON port swapping feature causes all configuration associated with *old-port-number* and *new port-number* to be swapped, including VSAN configurations.

Cisco MDS switches allow port swapping for nonexistent ports as follows:

- Only FICON-specific configurations (prohibit, block, and port address mapping) are swapped.
- No other system configuration is swapped.
- All other system configurations are only maintained for existing ports.

Tip If you check the **Active=Saved** check box on any FICON VSAN, then the swapped configuration is automatically saved to startup. Otherwise, you must explicitly save the running configuration immediately after swapping the ports.

Once you swap ports, the switch automatically performs the following actions:

- Shuts down both the old and new ports.
- Swaps the port configuration.

If you attempt to bring the port up, you must explicitly shut down the port to resume traffic.

Note To view the latest FICON information, you must click the **Refresh** button. See the "FICON Information Refresh Note" section on page 31-23.

This section includes the following topics:

About Swapping Ports

Be sure to follow these guidelines when using the FICON port swapping feature:

- Port swapping is not supported for logical ports (PortChannels, FCIP links). Neither the *old-port-number* nor the *new-port-number* can be a logical port.
- Port swapping is not supported between physical ports that are part of a PortChannel. Neither the *old-port-number* nor the *new-port-number* can be a physical port that is part of a PortChannel.
- Before performing a port swap, the Cisco SAN-OS software performs a compatibility check. If the two ports have incompatible configurations, the port swap is rejected with an appropriate reason code. For example, if a port with BB_credits as 25 is being swapped with an OSM port for which a maximum of 12 BB_credits is allowed (not a configurable parameter), the port swapping operation is rejected.
- Before performing a port swap, the Cisco SAN-OS software performs a compatibility check to verify the extended BB_credits configuration.
- If ports have default values (for some incompatible parameters), then a port swap operation is allowed and the ports retain their default values.
- Port tracking information is not included in port swapping. This information must be configured separately (see Chapter 61, "Configuring Port Tracking").

Note The 32-port module guidelines also apply for port swapping configurations.

Swapping Ports

To swap ports using Device Manager, follow these steps:

Step 1 Select two Fibre Channel ports by holding down the **CTRL** key and clicking them.

Step 2 Select **Swap Selected Ports** from the FICON menu (see Figure 31-21).

Figure 31-21 Swapping Ports with Device Manager

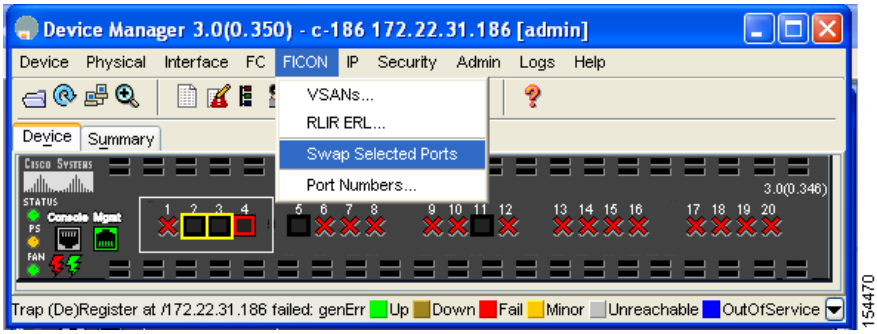

CUP In-Band Management

The Control Unit Port (CUP) protocol configures access control and provides unified storage management capabilities from a mainframe computer. Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for in-band management using the IBM S/A OS/390 I/O operations console.

Note The CUP specification is proprietary to IBM.

CUP is supported by switches and directors in the Cisco MDS 9000 Family. The CUP function allows the mainframe to manage the Cisco MDS switches.

Host communication includes control functions such as blocking and unblocking ports, as well as monitoring and error reporting functions.

This section includes the following topics:

Placing CUPs in a Zone

To place the CUP in a zone, follow these steps:

Step 1 In Fabric Manager, choose **Zone > Edit Full Zoneset,** and then choose **Edit > Edit Default Zone Attributes** to set the default zone to permit for the required VSAN. (See Figure 31-22.)

Figure 31-22 Setting the Default Zone Policy

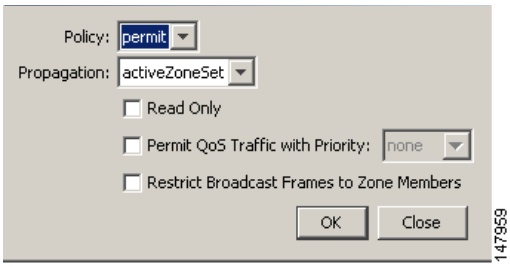

Step 2 In Device Manager, choose **FC > Name Server...** for the required VSAN and obtain the FICON:CUP WWN. See Figure 31-23.

Figure 31-23 Finding pWWN for FICON:CUP

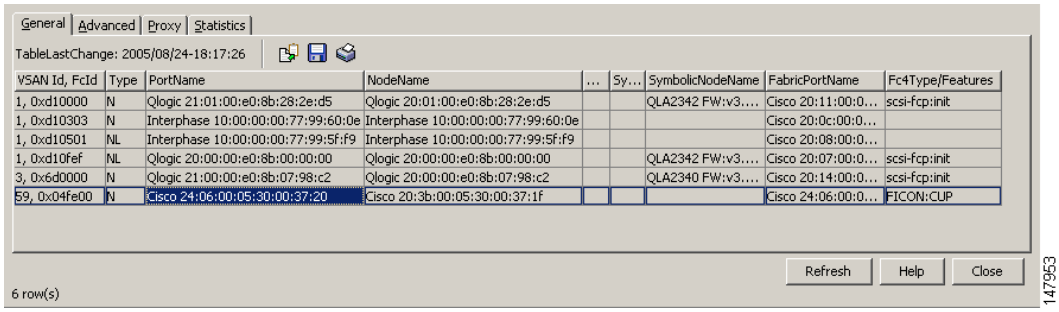

Note If more than one `FICON:CUP WWN` exists in this fabric, be sure to add all the `FICON:CUP` pWWNs to the required zone.

Step 3 In Fabric Manager, choose **Zone > Edit Full Zoneset** and add the FICON:CUP pWWN to the zone database. (See Figure 31-24.)

Figure 31-24 Adding FICON:CUP WWN to Zone

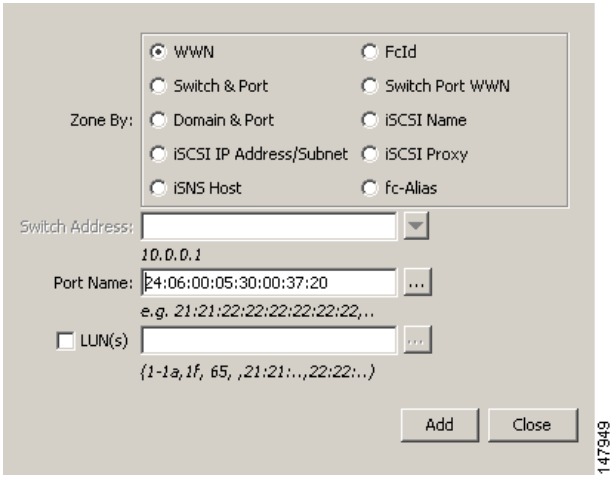

Receiving FICON Alerts

To receive an alert to indicate any changes in the FICON configuration using Device Manager, follow these steps:

Step 1 Select **VSANs** from the FICON menu.

You see the FICON VSAN configuration dialog box. The **VSANs** tab is default.

Step 2 Check the **User Alert Mode** check box to receive an alert when the FICON configuration changes.

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

Displaying FICON Port Address Information

To display FICON port address information using Device Manager, follow these steps:

Step 1 Select **VSANs** from the FICON menu.

You see the FICON VSAN configuration dialog box. The **VSANs** tab is default.

Step 2 Select a VSAN ID and click **Port Configuration**.

You see the FICON Port Configuration dialog box.

Step 3 Click **Close** to close the dialog box.

Displaying IPL File Information

To display the IPL file information using Device Manager, follow these steps:

Step 1 Select **VSANs** from the FICON menu.

You see the FICON VSAN configuration dialog box. The **VSANs** tab is default.

Step 2 Click the **Files** tab.

You see the FICON Files dialog box.

Step 3 Select the file that you want to view and click **Open**.

Step 4 Click **Close** to close the dialog box.

About the History Buffer

In the directory history buffer, the `Key Counter` column displays the 32-bit value maintained by Cisco MDS switches. This value is incremented when any port changes state in that VSAN. The key counter (a 32-bit value) is incremented when a FICON-related configuration is changed. Host programs can increment this value at the start of the channel program and then perform operations on multiple ports. The director history buffer keeps a log of which port address configuration was changed for each key-counter value.

The director history buffer provides a mechanism to determine the change in the port state from the previous time when a value was contained in the key counter.

Viewing the History Buffer

To view the directory history buffer using Device Manager, follow these steps:

Step 1 Select **VSANs** from the FICON menu.

You see the FICON VSAN configuration dialog box. The **VSANs** tab is default.

Step 2 Click the Director History button.

You see the history buffer dialog box.

Step 3 Click **Close** to close the dialog box.

Calculating FICON Flow Load Balance

The FICON Flow Load Balance Calculator allows you to get the best load balancing configuration for your FICON flows. The calculator does not rely on any switch or flow discovery in the fabric. It is available from the Fabric Manager Tools menu.

To use the FICON Flow Load Balance Calculator from Fabric Manager follow these steps:

Step 1 Click **Tools > Other > FICON Flow Load Balance Calculator**.

You see the Flow Load Balance Calculator (see Figure 31-25).

Figure 31-25 Flow Load Balance Calculator

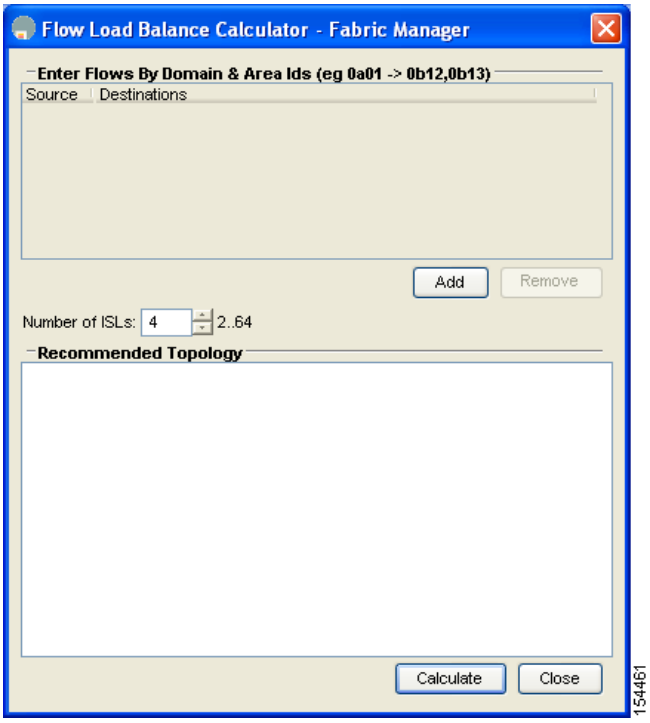

Step 2 Click **Add** to enter the source and destination(s) flows.

Step 3 Enter source and destination using 2 byte hex (by domain and area IDs).You can copy and paste these IDs, and then edit them if you need to (see Figure 31-25).

Figure 31-26 Flow Load Balance Calculator - Initial Screen

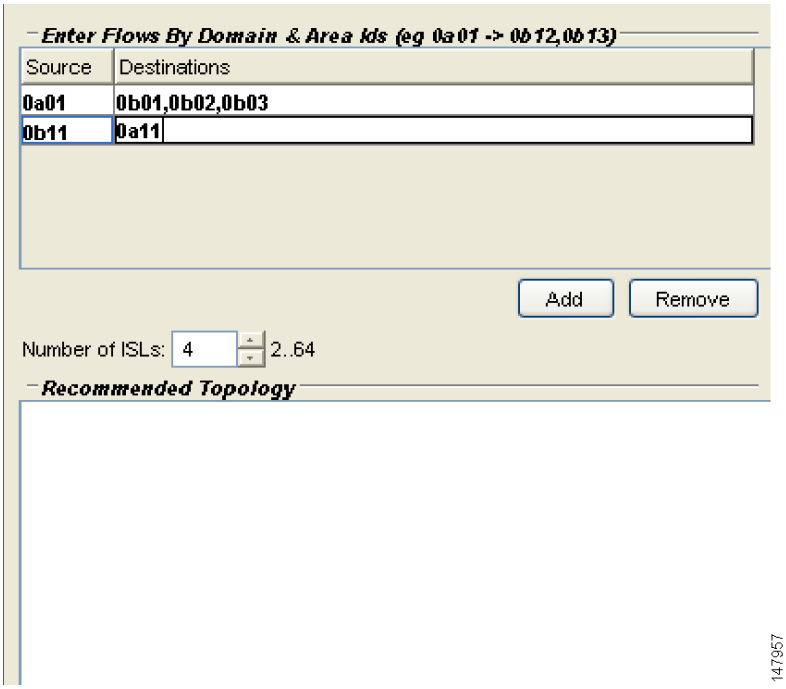

Step 4 Enter (or select) the number of ISLs between the two switches (for example, between domain ID 0a and 0b).

Step 5 Select a row to remove it and click **Remove**.

Step 6 Click **Calculate** to show the recommended topology.

Note If you change flows or ISLs, you must click **Calculate** to see the new recommendation.

Default Settings

Table 31-3 lists the default settings for FICON features.

Table 31-3 Default FICON Settings

Parameters	Default
FICON feature	Disabled.
Port numbers	Same as port addresses.
FC ID last byte value	0 (zero).
EBCDIC format option	US-Canada.
Switch offline state	Hosts can move the switch to an offline state.
Mainframe users	Users can configure FICON parameters on Cisco MDS switches.

Table 31-3 Default FICON Settings (continued)

Parameters	Default
Clock in each VSAN	Same as the switch hardware clock.
Host clock control	Host can set the clock on this switch.
SNMP users	Users can configure FICON parameters.
Port address	Not blocked.
Prohibited ports	Ports 90–253 and 255 for the Cisco MDS 9200 Series switches. Ports 250–253 and 255 for the Cisco MDS 9500 Series switches.

Table 31-4 lists the default settings for fabric binding features.

Table 31-4 Default Fabric Binding Settings

Parameters	Default
Fabric binding	Disabled.

CHAPTER 32

Advanced Features and Concepts

This chapter describes the advanced features provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

Common Interface Configuration

Common Information Model (CIM) is an object-oriented information model that extends the existing standards for describing management information in a network/enterprise environment. CIM messages are independent of platform and implementation because they are encoded in N Extensible Markup Language (XML). CIM consists of a specification and a schema. The specification defines the syntax and rules for describing management data and integrating with other management models. The schema provides the actual model descriptions for systems, applications, networks, and devices.

For more information about CIM, refer to the specification available through the Distributed Management Task Force (DMTF) website at the following URL: http://www.dmtf.org/

For further information about Cisco MDS 9000 Family support for CIM servers, refer to the *Cisco MDS 9000 Family CIM Programming Reference Guide*.

A CIM client is required to access the CIM server. The client can be any client that supports CIM.

For added security, you can install an SSL certificate to encrypt the logon information and enable the HTTPS server before enabling the CIM server. The CIM server is disabled by default. If you do not enable the HTTPS server, the standard HTTP server is enabled (default).

To configure a CIM server using the HTTPS or HTTP protocols, refer to the *Cisco MDS 9000 Family Configuration Guide*.

Some configuration settings are similar for Fibre Channel, management, and VSAN interfaces. You can configure interfaces from Fabric Manager by expanding **Switches > Interfaces** and selecting the interface type from the Physical Attributes pane. Figure 32-1 shows a sample of what you might see in the Information pane for Fibre Channel Interfaces.

Figure 32-1 Fibre Channel Interface Configuration

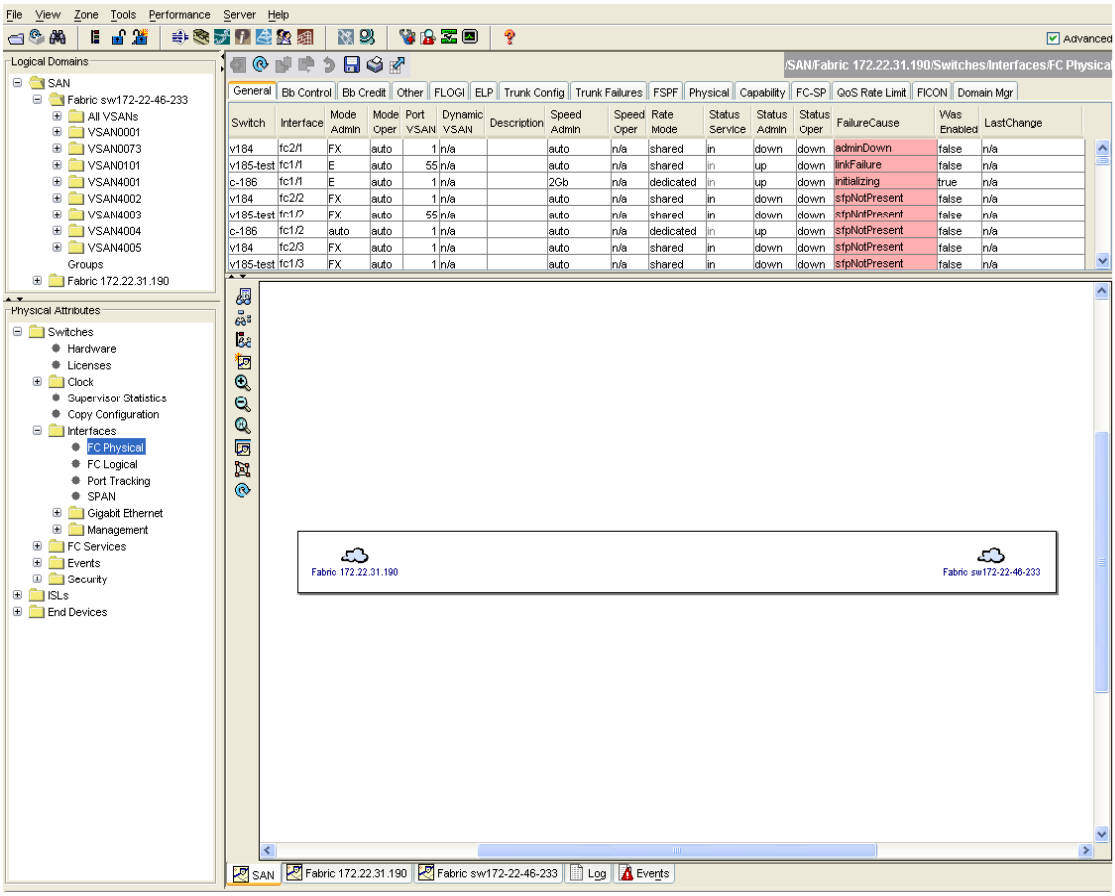

This section describes the common interface characteristics, including (but not limited to) states, speeds, and statistics. It includes the following topics:

About Interface States

The interface state depends on the administrative configuration of the interface and the dynamic state of the physical link.

Administrative States

The administrative state refers to the administrative configuration of the interface as described in Table 32-1.

Table 32-1 Administrative States

Administrative State	Description
Up	Interface is enabled.
Down	Interface is disabled. If you administratively disable an interface by shutting down that interface, the physical link layer state change is ignored.

Operational States

The operational state indicates the current operational state of the interface as described in Table 32-2.

Table 32-2 Operational States

Operational State	Description
Up	Interface is transmitting or receiving traffic as desired. To be in this state, an interface must be administratively up, the interface link layer state must be up, and the interface initialization must be completed.
Down	Interface cannot transmit or receive (data) traffic.
Trunking	Interface is operational in TE mode.

Reason Codes

Reason codes are dependent on the operational state of the interface as described in Table 32-3.

Table 32-3 Reason Codes for Interface States

Administrative Configuration	Operational Status	Reason Code
Up	Up	None.

Table 32-3 Reason Codes for Interface States (continued)

Administrative Configuration	Operational Status	Reason Code
Down	Down	Administratively down—If you administratively configure an interface as down, you disable the interface. No traffic is received or transmitted.
Up	Down	Administratively up and operationally down—The reason code differs based on the nonoperational reason code. See Table B-1 for a full description of these codes.

32-Port Configuration Guidelines

The 32-port guidelines apply to the following hardware:

- The 32-port 2-Gbps or 1-Gbps switching module
- The Cisco MDS 9140 Switch

When configuring these host-optimized ports, the following port mode guidelines apply:

- You can configure only the first port in each 4-port group (for example, the first port in ports 1-4, the fifth port in ports 5-8 and so on) as an E port. If the first port in the group is configured as an E port, the other three ports in each group (ports 2-4, 6-8 and so on) are not usable and remain shut down.
- If any of the other three ports are enabled, you cannot configure the first port as an E port. The other three ports continue to remain enabled.
- The auto mode is the default port mode. The auto mode is not allowed in a 32-port switching module or the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).
- The default port mode is Fx (Fx negotiates to F or FL) for 32-port switching modules and the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).
- The 32-port switching module does not support FICON.

Note In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the 9120 switch and 8 ports in the 9140 switch) are full line rate like the 16-port switching module. The other ports (16 ports in the 9120 switch and 32 ports in the 9140 switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.

Graceful Shutdown

Interfaces on a port are shut down, or disabled, by default (unless you modified the initial configuration).

The Cisco SAN-OS software implicitly performs a graceful shutdown in response to either of the following actions for interfaces operating in the E port mode:

- If you shut down an interface
- If a Cisco SAN-OS software application executes a port shut down as part of its function

A graceful shutdown ensures that no frames are lost when the interface is shutting down. When a shut down is triggered either by you or the Cisco SAN-OS software, the switches connected to the shutdown link coordinate with each other to ensure that all frames in the ports are safely sent through the link before shutting down. This enhancement reduces the chance of frame loss.

A graceful shutdown is not possible in the following situations:

- If you physically remove the port from the switch.
- If in-order-delivery (IOD) is enabled (see "In-Order Delivery" section on page 28-18)
- If the Min_LS_interval interval is higher than 10 seconds (see"Default Settings" section on page 28-25)

This feature is only triggered if both switches at either end of this E port interface are MDS switches and are using Cisco SAN-OS Release 2.0(1b) or later.

Setting the Interface Administrative State

To disable or enable an interface using Fabric Manager, follow these steps:

Step 1 Either expand **Switches > Interfaces** and then select **Gigabit Ethernet** or expand **Switches > Interfaces** and then select **FC Physical**. You see the interface configuration in the Information pane.

Step 2 Click the **General** tab.

Step 3 Click **Admin**.

You see the drop-down box shown in Figure 32-2.

Figure 32-2 Changing the Administrative Status of a Switch

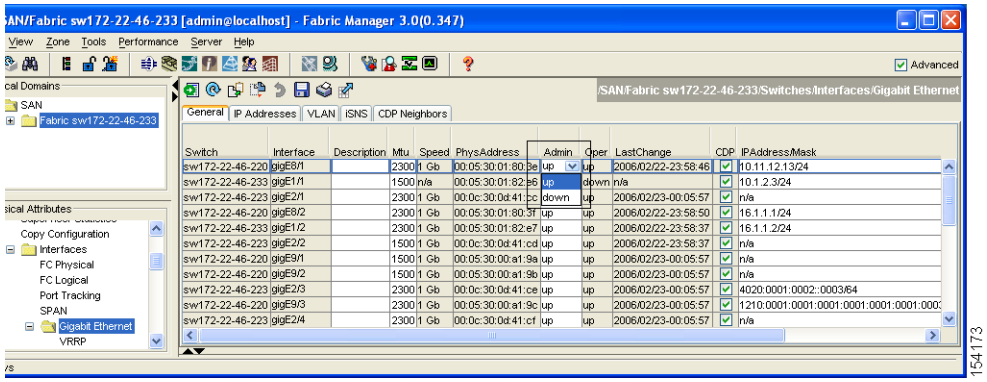

Step 4 Set the status to down (disable) or up (enable).

Step 5 Optionally, set other configuration parameters using the other tabs.

Step 6 Click **Apply Changes**.

About Administrative Speeds

By default, the administrative speed for an interface is automatically calculated by the switch.

Caution Changing the administrative speed is a disruptive operation.

Configuring the Administrative Speed

To configure the administrative speed of the interface using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **General** tab.

Step 3 Set the desired interface speed from the **Speed Admin** drop-down menu.Optionally, set other configuration parameters using the other tabs.

Step 4 Click **Apply Changes**.

About Interface Descriptions

Interface descriptions should help you identify the traffic or use for that interface. The interface description can be any alphanumeric string.

Configuring the Interface Description

To configure a description for an interface using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** then select **FC Physical**. You see the interface configuration in the Information pane.

Step 2 Choose the **General** tab.

Step 3 Click the Description field.

Step 4 Optionally, set other configuration parameters using the other tabs.

Step 5 Click **Apply Changes**.

About Beacon Mode

By default, the beacon mode is disabled on all switches. The beacon mode is indicated by a flashing green light that helps you identify the physical location of the specified interface.

Beacon mode has no effect on the operation of the interface.

Configuring Beacon Mode

To enable beacon mode for a specified interface or range of interfaces using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Interfaces** then select **FC Physical**. You see the interface configuration in the Information pane.

Step 2 Choose the **Physical** tab

You see the screen shown in Figure 32-3.

Figure 32-3 Physical Tab Contents for an FC Physical Interface

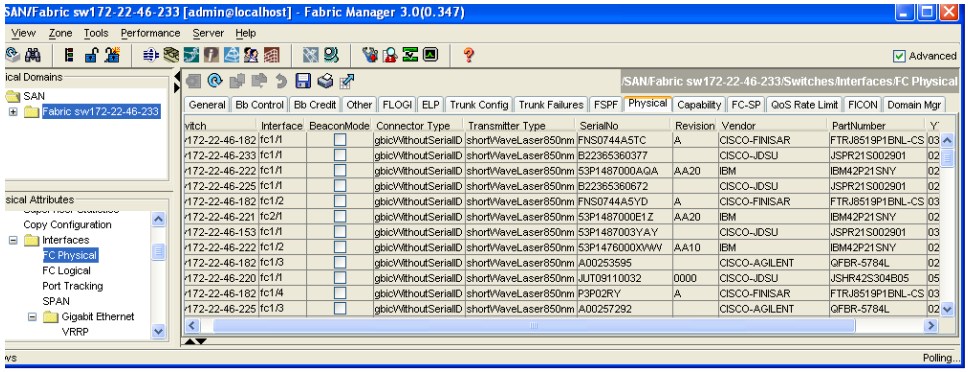

Step 3 Check the **BeaconMode** check box to enable beacon mode for that interface.

Step 4 Optionally, set other configuration parameters using the other tabs.

Step 5 Click **Apply Changes**.

Note The flashing green light turns on automatically when an external loopback is detected that causes the interfaces to be isolated. The flashing green light overrides the beacon mode configuration. The state of the LED is restored to reflect the beacon mode configuration after the external loopback is removed.

Identifying the LEDs

Figure 32-4 displays the status, link, and speed LEDs in a 16-port switching module.

Figure 32-4 Cisco MDS 9000 Family Switch Module LEDs

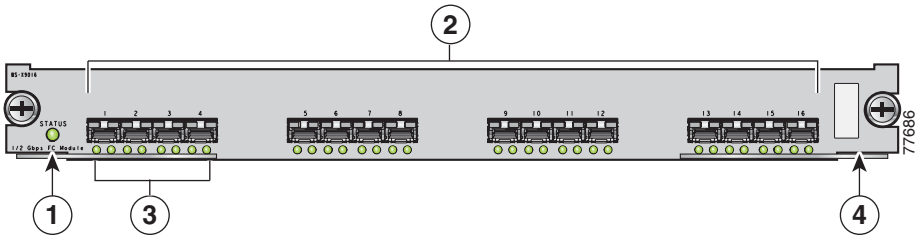

1	Status LED[1]	**3**	Link LEDs[1] and speed LEDs[2]
2	1/2-Gbps Fibre Channel port group[3]	**4**	Asset tag[4]

1. See the "Identifying Module LEDs" section on page 17-10.
2. See the "Identifying the LEDs" section on page 32-8.
3. See the "Graceful Shutdown" section on page 32-4.
4. See the Cisco MDS 9000 Family Hardware Installation Guides

Each port has one link LED on the left and one speed LED on the right.

The speed LED displays the speed of the port interface:

- Off—The interface attached to that port is functioning at 1000 Mbps.
- On (solid green)—The interface attached to that port is functioning at 2000 Mbps (for 2 Gbps interfaces).

The speed LED also displays if the beacon mode is enabled or disabled:

- Off—Beacon mode is disabled.
- On (flashing green)—The beacon mode is enabled. The LED flashes at one-second intervals.

About Attribute Default Values for Switch Ports

You can configure attribute default values for various switch port attributes. These attributes will be applied globally to all future switch port configurations, even if you do not individually specify them at that time.

Refer to the *Cisco MDS 9000 Family CLI Configuration Guide* to configure switch port attributes.

About Gathering Interface Statistics

You can use Device Manager to collect interface statistics on any switch. These statistics are collected at intervals that you can set.

Gathering Interface Statistics

To gather and display interface counters using Device Manager, follow these steps:

Step 1 Right-click an interface and select **Monitor**.

You see the Interface Monitor dialog box.

Step 2 Set both the number of seconds at which you want to poll the interface statistics and how you want the data represented in the Interval drop-down menus. For example, click **10s** and **LastValue/sec** as shown in Figure 18-3.

Step 3 Select any tab shown in Figure 32-5 to view those related statistics.

Figure 32-5 Device Manager Interface Monitor Dialog Box

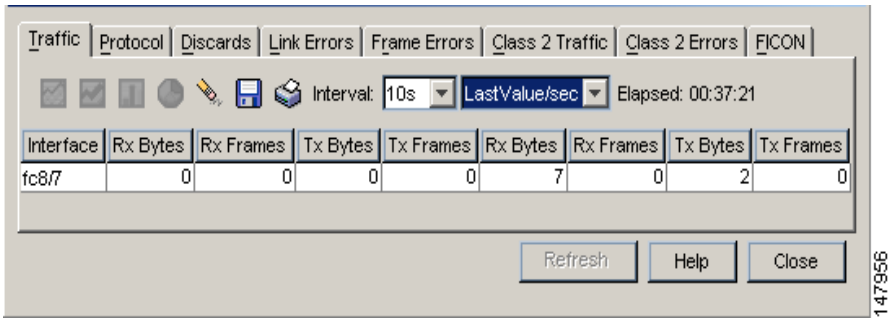

Step 4 Optionally, click the **Pencil** icon (the yellow eraser in Figure 32-5) to reset the cumulative counters.

Step 5 Optionally, click the **Save** icon to save the gathered statistics to a file or select the **Print** icon to print the statistics.

Step 6 Click **Close** when you are finished gathering and displaying statistics.

Fibre Channel Time Out Values

You can modify Fibre Channel protocol related timer values for the switch by configuring the following time out values (TOVs):

- Distributed services TOV (D_S_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 5,000 milliseconds.
- Error detect TOV (E_D_TOV)—The valid range is from 1,000 to 10,000 milliseconds. The default is 2,000 milliseconds. This value is matched with the other end during port initialization.
- Resource allocation TOV (R_A_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 10,000 milliseconds. This value is matched with the other end during port initialization.

Note The fabric stability TOV (F_S_TOV) constant cannot be configured.

This section includes the following topics:

Timer Configuration Across All VSANs

You can modify Fibre Channel protocol related timer values for the switch.

Caution The D_S_TOV, E_D_TOV, and R_A_ TOV values cannot be globally changed unless all VSANs in the switch are suspended.

Note If a VSAN is not specified when you change the timer value, the changed value is applied to all VSANs in the switch.

Configuring Timers

To configure timers in Fabric Manager, expand **Switches > FC Services** and then select **Timers & Policies** in the Physical Attributes pane. You see the timers for multiple switches in the Information pane. Click the **Change Timeouts** button to configure the timeout values. You see the dialog box in Figure 32-6.

Figure 32-6 Configure Timers in Fabric Manager

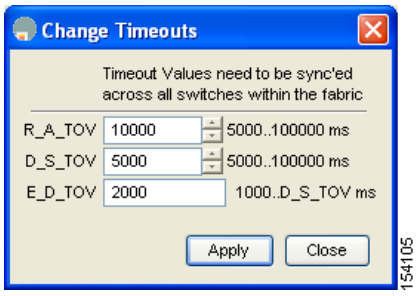

To configure timers in Device Manager, click **FC > Advanced > Timers/Policies**. You see the timers for a single switch in the dialog box shown in Figure 32-7.

Figure 32-7 Configure Timers in Device Manager

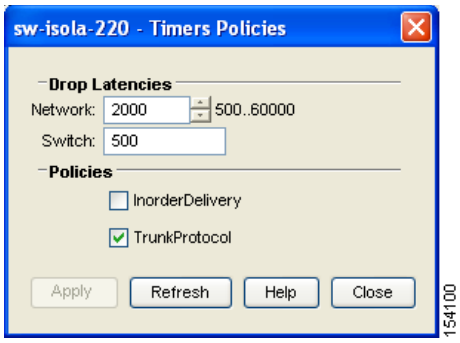

About Per-VSAN Timers

You can also issue the fctimer for a specified VSAN to configure different TOV values for VSANs with special links like FC or IP tunnels. You can configure different E_D_TOV, R_A_TOV, and D_S_TOV values for individual VSANs. Active VSANs are suspended and activated when their timer values are changed.

Caution You cannot perform a nondisruptive downgrade to any earlier version that does not support per-VSAN FC timers.

Note This configuration must be propagated to all switches in the fabric—be sure to configure the same value in all switches in the fabric.

If a switch is downgraded to Cisco MDS SAN-OS Release 1.2 or 1.1 after the timer is configured for a VSAN, an error message is issued to warn against strict incompatibilities. See the *Cisco MDS 9000 Family Troubleshooting Guide*.

Configuring Per-VSAN Timers

To configure per-VSAN FC timers using Device Manager, follow these steps:

Step 1 Click **FC > Advanced > VSAN Timers**.

You see the VSANs Timer dialog box shown in Figure 32-8.

Figure 32-8 VSAN Timers in Device Manager

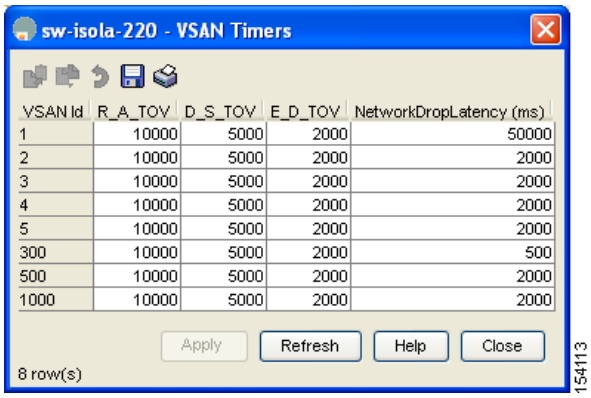

Step 2 Fill in the timer values that you want to configure.

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

About fctimer Distribution

You can enable per-VSAN fctimer fabric distribution for all Cisco MDS switches in the fabric. When you perform fctimer configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you issue the first configuration command after you enabled distribution in a switch. The fctimer application uses the effective and pending database model to store or commit the commands based on your configuration.

Refer to Chapter 12, "Using the CFS Infrastructure" for more information on the CFS application.

When you commit the fctimer configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. When you commit the fctimer configuration changes without implementing the session feature, the fctimer configurations are distributed to all the switches in the physical fabric.

Enabling or Disabling fctimer Distribution

To enable and distribute fctimer configuration changes using Device Manager, follow these steps:

Step 1 Choose **FC > Advanced > VSAN Timers**.

You see the VSANs Timer dialog box.

Figure 32-9 VSAN Timers in Device Manager

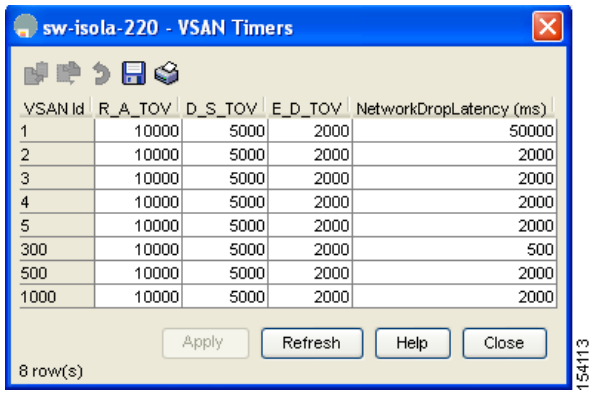

Step 2 Fill in the timer values that you want to configure.

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

Step 4 Select **commit** from the CFS drop-down menu to distribute these changes or select **abort** from the CFS drop-down menu to discard any unsaved changes.

Database Merge Guidelines

See the "CFS Merge Support" section on page 12-9 for detailed concepts.

When merging two fabrics, follow these guidelines:

- Be aware of the following merge conditions:
  - The merge protocol is not implemented for distribution of the fctimer values—you must manually merge the fctimer values when a fabric is merged.The per-VSAN fctimer configuration is distributed in the physical fabric.
  - The fctimer configuration is only applied to those switches containing the VSAN with a modified fctimer value.
  - The global fctimer values are not distributed.
- Do not configure global timer values when distribution is enabled.

Note The number of pending fctimer configuration operations cannot be more than 15. At that point, you must commit or abort the pending configurations before performing any more operations.

World Wide Names

The world wide name (WWN) in the switch is equivalent to the Ethernet MAC address. As with the MAC address, you must uniquely associate the WWN to a single device. The principal switch selection and the allocation of domain IDs rely on the WWN. The WWN manager, a process-level manager residing on the switch's supervisor module, assigns WWNs to each switch.

Cisco MDS 9000 Family switches support three network address authority (NAA) address formats (see Table 32-4).

Table 32-4 Standardized NAA WWN Formats

NAA Address	NAA Type	WWN Format	
IEEE 48-bit address	Type 1 = 0001b	000 0000 0000b	48-bit MAC address
IEEE extended	Type 2 = 0010b	Locally assigned	48-bit MAC address
IEEE registered	Type 5 = 0101b	IEEE company ID: 24 bits	VSID: 36 bits

Caution Changes to the world-wide names should be made by an administrator or individual who is completely familiar with switch operations.

This section includes the following topics:

Displaying WWN Information

To display WWN information using Device Manager, choose **FC > Advanced > WWN Manager.** You see the list of allocated WWNs.

Link Initialization WWN Usage

Exchange Link Protocol (ELP) and Exchange Fabric Protocol (EFP) use WWNs during link initialization. The usage details differ based on the Cisco SAN-OS software release:

Both ELPs and EFPs use the VSAN WWN by default during link initialization. However, the ELP usage changes based on the peer switch's usage:

- If the peer switch ELP uses the switch WWN, then the local switch also uses the switch WWN.
- If the peer switch ELP uses the VSAN WWN, then the local switch also uses the VSAN WWN.

Note As of Cisco SAN-OS Release 2.0(2b), the ELP is enhanced to be compliant with FC-SW-3.

Configuring a Secondary MAC Address

To allocate secondary MAC addresses using Device Manager, follow these steps:

Step 1 Choose **FC > Advanced > WWN Manager.**

You see the list of allocated WWNs shown in Figure 32-10.

Figure 32-10 Allocated World Wide Names in Device Manager

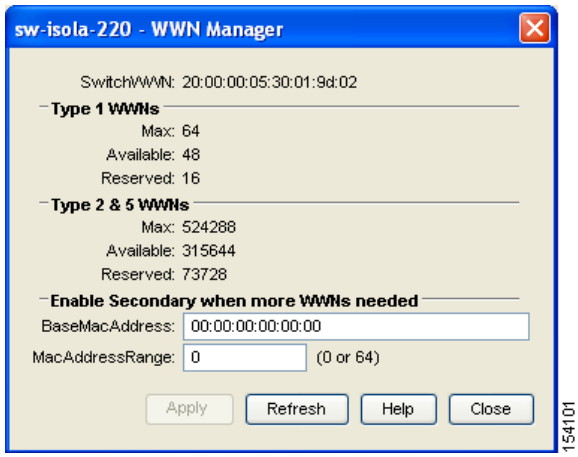

Step 2 Supply the BaseMacAddress and MacAddressRange fields

Step 3 Click **Apply** to save these changes or click **Close** to discard any unsaved changes.

FC ID Allocation for HBAs

Fibre Channel standards require a unique FC ID to be allocated to an N port attached to a Fx port in any switch. To conserve the number of FC IDs used, Cisco MDS 9000 Family switches use a special allocation scheme.

Some HBAs do not discover targets that have FC IDs with the same domain and area. Prior to Cisco SAN-OS Release 2.0(1b), the Cisco SAN-OS software maintained a list of tested company IDs that do not exhibit this behavior. These HBAs were allocated with single FC IDs, and for others a full area was allocated.

The FC ID allocation scheme available in Release 1.3 and earlier, allocates a full area to these HBAs. This allocation isolates them to that area and are listed with their pWWN during a fabric login. The allocated FC IDs are cached persistently and are still available in Cisco SAN-OS Release 2.0(1b) (see the "FC ID Allocation for HBAs" section on page 32-15).

To allow further scalability for switches with numerous ports, the Cisco SAN-OS software maintains a list of HBAs exhibiting this behavior. Each HBA is identified by its company ID (also known as Organizational Unique Identifier, or OUI) used in the pWWN during a fabric log in. Hence a full area is allocated to the N ports with company IDs that are listed and for the others, a single FC ID is allocated. Irrespective of the kind (whole area or single) of FC ID allocated, the FC ID entries remain persistent.

This section includes the following topics:

Default Company ID list

All switches in the Cisco MDS 9000 Family that ship with Cisco SAN-OS Release 2.0(1b) or later, contain a default list of company IDs that require area allocation. Using the company ID reduces the number of configured persistent FC ID entries. You can configure or modify these entries using the CLI.

Note Persistent entries take precedence over company ID configuration. If the HBA fails to discover a target, verify that the HBA and the target are connected to the same switch and have the same area in their FC IDs, then perform the following procedure:

1. Shut down the port connected to the HBA.
2. Clear the persistent FC ID entry.
3. Get the company ID from the Port WWN.
4. Add the company ID to the list that requires area allocation.
5. Bring up the port.

The list of company IDs have the following characteristics:

- A persistent FC ID configuration always takes precedence over the list of company IDs. Hence even if the company ID is configured to receive an area, the persistent FC ID configuration results in the allocation of a single FC ID.
- New company IDs added to subsequent releases are automatically added to existing company IDs.
- The list of company IDs is saved as part of the running and saved configuration.
- The list of company IDs is used only when the fcinterop FC ID allocation scheme is in auto mode. By default, the interop FC ID allocation is set to auto, unless changed.

Tip We recommend that you set the fcinterop FC ID allocation scheme to auto and use the company ID list and persistent FC ID configuration to manipulate the FC ID device allocation.

Refer to the *Cisco MDS 9000 Family CLI Configuration Guide* to change the FC ID allocation.

Verifying Company ID Configuration

To view the configured company IDs using Device Manager, choose **FC > Advanced > FcId Area Allocation**. You can implicitly derive the default entries shipped with a specific release by combining the list of Company IDs displayed without any identification with the list of deleted entries.

Some WWN formats do not support company IDs. In these cases, you may need to configure the FC ID persistent entry.

Switch Interoperability

Interoperability enables the products of multiple vendors to come into contact with each other. Fibre Channel standards guide vendors towards common external Fibre Channel interfaces.

If all vendors followed the standards in the same manner, then interconnecting different products would become a trivial exercise. However, not all vendors follow the standards in the same way, thus resulting in interoperability modes. This section briefly explains the basic concepts of these modes.

Each vendor has a regular mode and an equivalent interoperability mode, which specifically turns off advanced or proprietary features and provides the product with a more amiable standards compliant implementation.

This section includes the following topics:

About Interop Mode

Cisco SAN-OS software supports the following four interop modes:

- Mode 1— Standards based interop mode that requires all other vendors in the fabric to be in interop mode.
- Mode 2—Brocade native mode (Core PID 0).
- Mode 3—Brocade native mode (Core PID 1).
- Mode 4—McData native mode.

For information about configuring interop modes 2, 3, and 4, refer to the *MDS Switch to Switch Interoperability Configuration Guide*.

Table 32-5 lists the changes in switch behavior when you enable interoperability mode. These changes are specific to switches in the Cisco MDS 9000 Family while in interop mode.

Table 32-5 Changes in Switch Behavior When Interoperability Is Enabled

Switch Feature	Changes if Interoperability Is Enabled
Domain IDs	Some vendors cannot use the full range of 239 domains within a fabric. Domain IDs are restricted to the range 97-127. This is to accommodate McData's nominal restriction to this same range. They can either be set up statically (the Cisco MDS switch accept only one domain ID, if it does not get that domain ID it isolates itself from the fabric) or preferred. (If it does not get its requested domain ID, it accepts any assigned domain ID.)
Timers	All Fibre Channel timers must be the same on all switches as these values are exchanged by E ports when establishing an ISL. The timers are F_S_TOV, D_S_TOV, E_D_TOV, and R_A_TOV.
F_S_TOV	Verify that the Fabric Stability Time Out Value timers match exactly.
D_S_TOV	Verify that the Distributed Services Time Out Value timers match exactly.
E_D_TOV	Verify that the Error Detect Time Out Value timers match exactly.

Table 32-5 Changes in Switch Behavior When Interoperability Is Enabled (continued)

Switch Feature	Changes if Interoperability Is Enabled
R_A_TOV	Verify that the Resource Allocation Time Out Value timers match exactly.
Trunking	Trunking is not supported between two different vendor's switches. This feature may be disabled on a per port or per switch basis.
Default zone	The default zone behavior of permit (all nodes can see all other nodes) or deny (all nodes are isolated when not explicitly placed in a zone) may change.
Zoning attributes	Zones may be limited to the pWWN and other proprietary zoning methods (physical port number) may be eliminated. **Note** Brocade uses the **cfgsave** command to save fabric-wide zoning configuration. This command does not have any effect on Cisco MDS 9000 Family switches if they are part of the same fabric. You must explicitly save the configuration on each switch in the Cisco MDS 9000 Family.
Zone propagation	Some vendors do not pass the full zone configuration to other switches, only the active zone set gets passed. Verify that the active zone set or zone configuration has correctly propagated to the other switches in the fabric.
VSAN	Interop mode only affects the specified VSAN.
TE ports and PortChannels	TE ports and PortChannels cannot be used to connect Cisco MDS to non-Cisco MDS switches. Only E ports can be used to connect to non-Cisco MDS switches. TE ports and PortChannels can still be used to connect an Cisco MDS to other Cisco MDS switches even when in interop mode.
FSPF	The routing of frames within the fabric is not changed by the introduction of interop mode. The switch continues to use src-id, dst-id, and ox-id to load balance across multiple ISL links.
Domain reconfiguration disruptive	This is a switch-wide impacting event. Brocade and McData require the entire switch to be placed in offline mode and/or rebooted when changing domain IDs.
Domain reconfiguration nondisruptive	This event is limited to the affected VSAN. Only Cisco MDS 9000 Family switches have this capability—only the domain manager process for the affected VSAN is restarted and not the entire switch.
Name server	Verify that all vendors have the correct values in their respective name server database.
IVR	IVR-enabled VSANs can be configured in **no interop** (default) mode or in any of the **interop** modes.

Configuring Interoperability

The interop mode in Cisco MDS 9000 Family switches can be enabled disruptively or nondisruptively.

Note Brocade's **msplmgmtdeactivate** command must explicitly be run prior to connecting from a Brocade switch to either Cisco MDS 9000 Family switches or to McData switches. This command uses Brocade proprietary frames to exchange platform information, which Cisco MDS 9000 Family switches or McData switches do not understand. Rejecting these frames causes the common E ports to become isolated.

To configure the interop mode for a VSAN using Fabric Manager, follow these steps:

Step 1 Choose **VSAN*xxx* > VSAN Attributes** from the Logical Domains pane.

You see the VSAN attributes in the Information pane shown in Figure 32-11.

Figure 32-11 VSAN Attributes in Fabric Manager

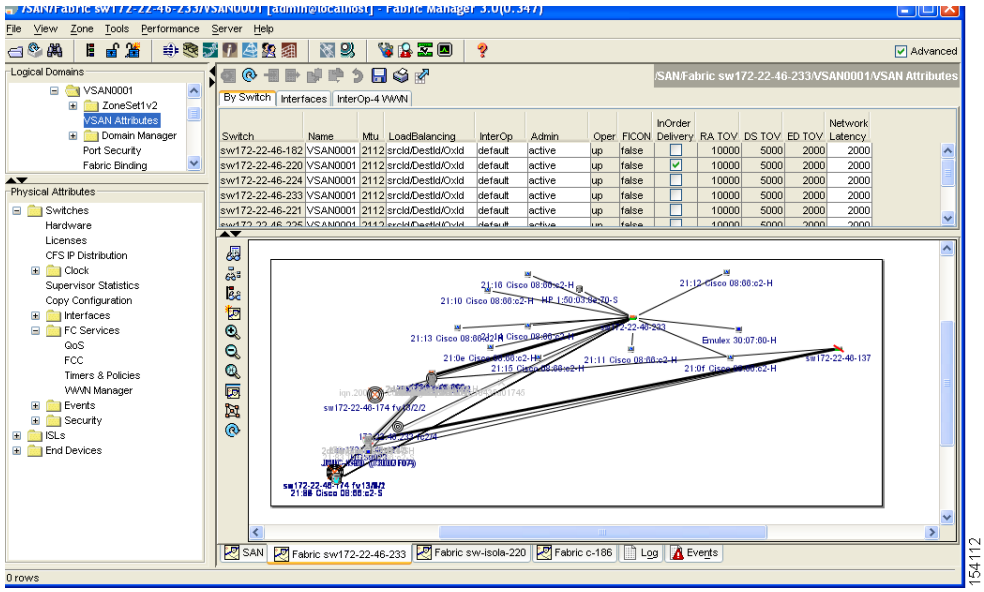

Step 2 Select **Interop-1** from the InterOp drop-down menu as shown in Figure 32-12.

Figure 32-12 Interop Drop-down Menu

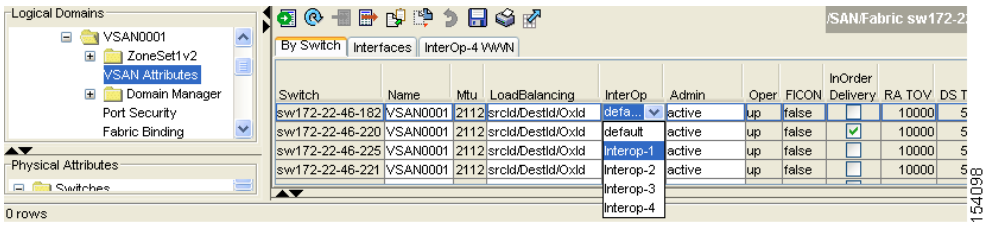

Step 3 Click **Apply Changes** to save this interop mode.

Step 4 Expand **VSAN*xxx*** and then select **Domain Manager** from the Logical Domains pane.

You see the Domain Manager configuration in the Information pane as shown in Figure 32-13.

Figure 32-13 Domain Manager Configuration

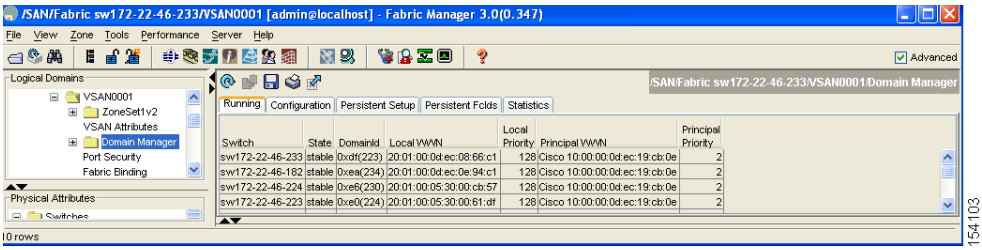

Step 5 Set the Domain ID in the range of 97 (0x61) through 127 (0x7F).

a. Click the **Configuration** tab (see Figure 32-13).

b. Click in the Configure Domain ID column under the Configuration tab.

c. Click the **Running** tab and check that the change has been made.

Note This is a limitation imposed by the McData switches.

Note When changing the domain ID, the FC IDs assigned to N ports also change.

Step 6 Change the Fibre Channel timers (if they have been changed from the system defaults).

Note The Cisco MDS 9000, Brocade, and McData FC error detect (ED_TOV) and resource allocation (RA_TOV) timers default to the same values. They can be changed if needed. The RA_TOV default is 10 seconds, and the ED_TOV default is 2 seconds. Per the FC-SW2 standard, these values must be the same on each switch within the fabric.

a. Expand **Switches > FC Services** and then select **Timers and Policies**. You see the timer settings in the Information pane.

b. Click **Change Timeouts** to modify the time-out values. You see the Change Timeouts Dialog Box in Figure 32-14.

Figure 32-14 Change TImeouts Dialog Box

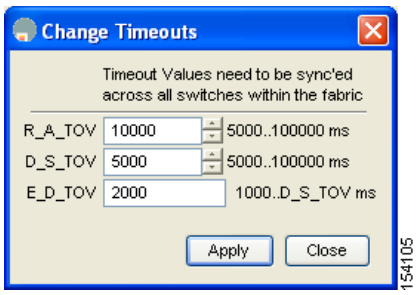

c. Click **Apply** to save the new time-out values.

Step 7 Optionally, choose **VSAN*xxx* > Domain Manager> Configuration** tab and select **disruptive** or **nonDisruptive** in the Restart column to restart the domain.

Verifying Interoperating Status

This section highlights the steps used to verify if the fabric is up and running in interoperability mode.

To verify the interoperability status of any switch in the Cisco MDS 9000 Family using Fabric Manager, follow these steps:

Step 1 Choose **Switches** in the Physical Attributes pane and check the release number in the Information pane to verify the Cisco SAN-OS release.

Step 2 Expand **Switches > Interfaces** and then select **FC Physical** to verify the interface modes for each switch.

You see the **General** tab contents shown in Figure 32-15. Interface is the second column.

Figure 32-15 Interface Modes for a Switch

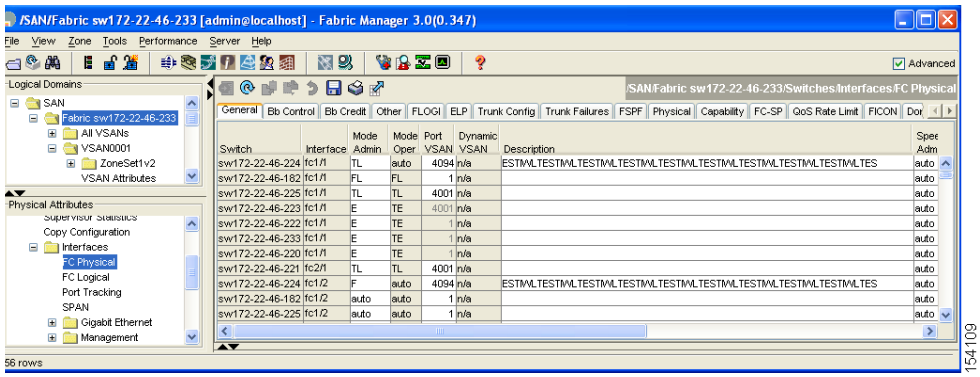

Step 3 Expand **Fabric*xx*** in the Logical Domains pane and then select **All VSANs** to verify the interop mode for all VSANs.

Step 4 Expand **Fabric*xx* > All VSANs** and then select **Domain Manager** to verify the domain IDs, local, and principal sWWNs for all VSANs. See Figure 32-16.

Figure 32-16 Domain Manager Information

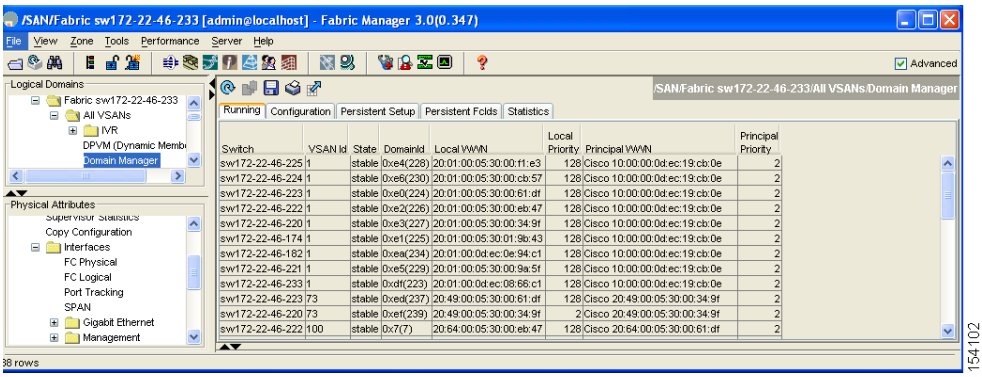

Step 5 Using Device Manager, choose **FC > Name Server** to verify the name server information. See Figure 32-17

Figure 32-17 Name Server in Device Manager

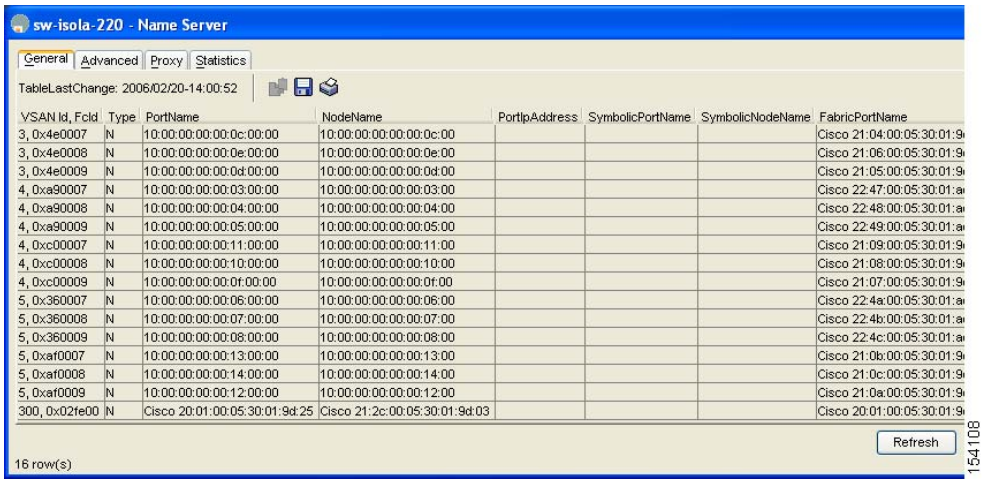

Note The Cisco MDS name server shows both local and remote entries, and does not time out the entries.

Default Settings

Table 32-6 lists the default settings for the features included in this chapter.

Table 32-6 Default Settings for Advanced Features

Parameters	Default
CIM server	Disabled
CIM server security protocol	HTTP
D_S_TOV	5,000 milliseconds.
E_D_TOV	2,000 milliseconds.
R_A_TOV	10,000 milliseconds.
Timeout period to invoke fctrace	5 seconds.
Number of frame sent by the fcping feature	5 frames.
Remote capture connection protocol	TCP.
Remote capture connection mode	Passive.
Local capture frame limit s	10 frames.
FC ID allocation mode	Auto mode.
Loop monitoring	Disabled.

Cisco Systems

Send documentation comments to mdsfeedback-doc@cisco.com

PART 5

Security

CHAPTER 33

Configuring Users and Common Roles

The CLI and SNMP use common roles in all switches in the Cisco MDS 9000 Family. You can use the CLI to modify a role that was created using SNMP and vice versa.

Users, passwords, and roles for all CLI and SNMP users are the same. A user configured through the CLI can access the switch using SNMP (for example, the Fabric Manager or the Device Manager) and vice versa.

This chapter includes the following sections:

Role-Based Authorization

Switches in the Cisco MDS 9000 Family perform authentication based on roles. Role-based authorization limits access to switch operations by assigning users to roles. This kind of authentication restricts you to management operations based on the roles to which you have been added.

When you execute a command, perform command completion, or obtain context sensitive help, the switch software allows the operation to progress if you have permission to access that command.

This section includes the following topics:

About Roles

Each role can contain multiple users and each user can be part of multiple roles. For example, if role1 users are only allowed access to configuration commands, and role2 users are only allowed access to debug commands, then if Joe belongs to both role1 and role2, he can access configuration as well as debug commands.

Note If you belong to multiple roles, you can execute a union of all the commands permitted by these roles. Access to a command takes priority over being denied access to a command. For example, suppose you belong to a TechDocs group and you were denied access to configuration commands. However, you also belong to the engineering group and have access to configuration commands. In this case, you will have access to configuration commands.

Tip Any role, when created, does not allow access to the required commands immediately. The administrator must configure appropriate rules for each role to allow access to the required commands.

Configuring Roles and Profiles

To create an additional role or to modify the profile for an existing role using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane. Click the **Roles** tab in the Information pane.

You see the information Figure 33-1

Figure 33-1 Roles Tab in Users and Roles Screen

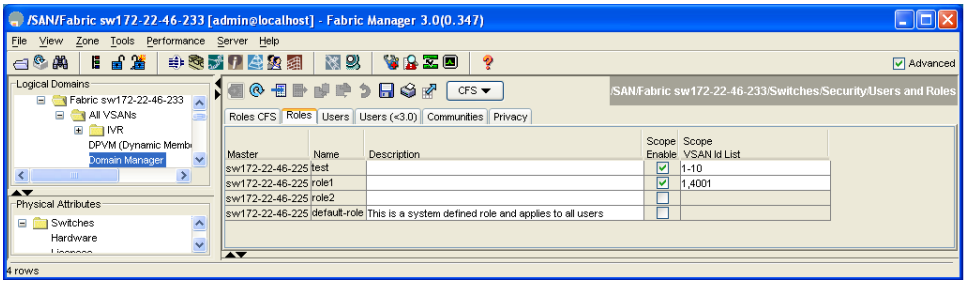

Step 2 Click **Create Row** to create a role in Fabric Manager.

You see the Roles - Create dialog box in Figure 33-2.

Figure 33-2 Create Roles Dialog Box

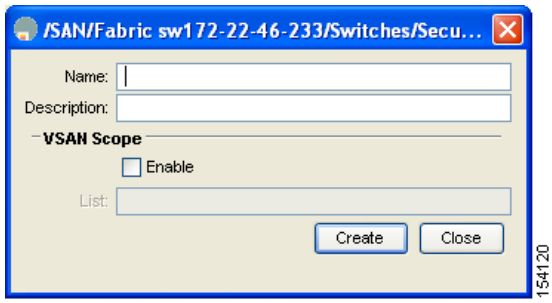

Step 3 Select the switches on which to configure a role.

Step 4 Enter the name of the role in the Name field.

Step 5 Enter the description of the role in the Description field.

Step 6 Optionally, check the **Enable** check box to enable the VSAN scope and enter the list of VSANs in the Scope field to which you want to restrict this role.

Step 7 Click **Create** to create the role, or click **Close** to close the Roles - Create dialog box without creating the common role.

Note Device Manager automatically creates six roles that are required for Device Manager to display a view of a switch. These roles are: **system**, **snmp**, **module**, **interface**, **hardware**, and **environment**.

Deleting Common Roles

To delete a common role using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane. Click the **Roles** tab in the Information pane.

Step 2 Click the role you want to delete.

Step 3 Click **Delete Row** to delete the common role.

Step 4 Click **Yes** to confirm the deletion or **No** to cancel it.

About the VSAN Policy

Configuring the VSAN policy requires the ENTERPRISE_PKG license (see Chapter 10, "Obtaining and Installing Licenses").

You can configure a role so that it only allows tasks to be performed for a selected set of VSANs. By default, the VSAN policy for any role is permit, which allows tasks to be performed for all VSANs. You can configure a role that only allows tasks to be performed for a selected set of VSANs. To selectively allow VSANs for a role, set the VSAN policy to deny, and then set the configuration to permit or the appropriate VSANs.

Note Users configured in roles where the VSAN policy is set to deny cannot modify the configuration for E ports. They can only modify the configuration for F or FL ports (depending on whether the configured rules allow such configuration to be made). This is to prevent such users from modifying configurations that may impact the core topology of the fabric.

Tip Roles can be used to create VSAN administrators. Depending on the configured rules, these VSAN administrators can configure MDS features (for example, zone, fcdomain, or VSAN properties) for their VSANs without affecting other VSANs. Also, if the role permits operations in multiple VSANs, then the VSAN administrators can change VSAN membership of F or FL ports among these VSANs.

Users belonging to roles in which the VSAN policy is set to deny are referred to as VSAN-restricted users.

Modifying the VSAN Policy

To modify the VSAN policy for an existing role using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane. Click the **Roles** tab in the Information pane.

Step 2 Check the **Scope Enable** check box if you want to enable the VSAN scope and restrict this role to a subset of VSANs.

Step 3 Enter the list of VSANs in the Scope VSAN Id List field that you want to restrict this role to.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

About Rules and Features for Each Role

Up to 16 rules can be configured for each role. These rules reflect what CLI commands are allowed. The user-specified rule number determines the order in which the rules are applied. For example, rule 1 is applied before rule 2, which is applied before rule 3, and so on. A user not belonging to the network-admin role cannot perform commands related to roles.

For example, if user A is permitted to perform all **show** CLI commands, user A cannot view the output of the **show role** CLI command if user A does not belong to the network-admin role

A rule specifies operations that can be performed by a specific role. Each rule consists of a rule number, a rule type (permit or deny), a CLI command type (for example, **config**, **clear**, **show**, **exec**, **debug**), and an optional feature name (for example, FSPF, zone, VSAN, fcping, or interface).

Note In this case, **exec** CLI commands refer to all commands in the EXEC mode that do not fall in the **show**, **debug**, and **clear**, CLI command categories.

Modifying Rules

To modify the rules for an existing role using Device Manager, follow these steps:

Step 1 Click **Security** > **Roles**.

Step 2 You see the Common Roles dialog box shown in Figure 33-3.

Figure 33-3 Common Roles Dialog Box in Device Manager

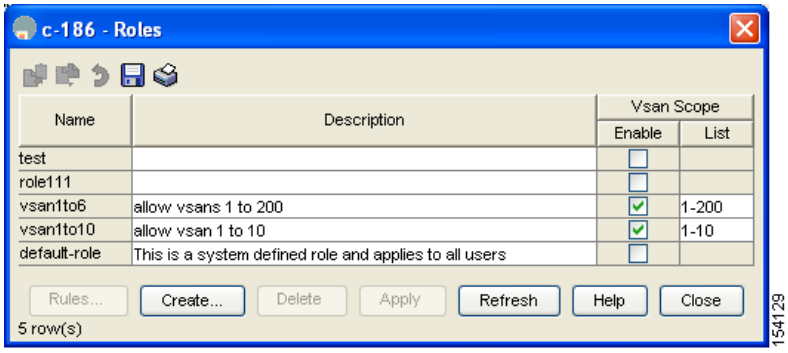

Step 3 Click the role for which you want to edit the rules.

Step 4 Click **Rules** to view the rules for the role.

You see the Rules dialog box shown in Figure 33-4. It may take a few minutes to display.

Figure 33-4 Edit Common Role Rules Dialog Box

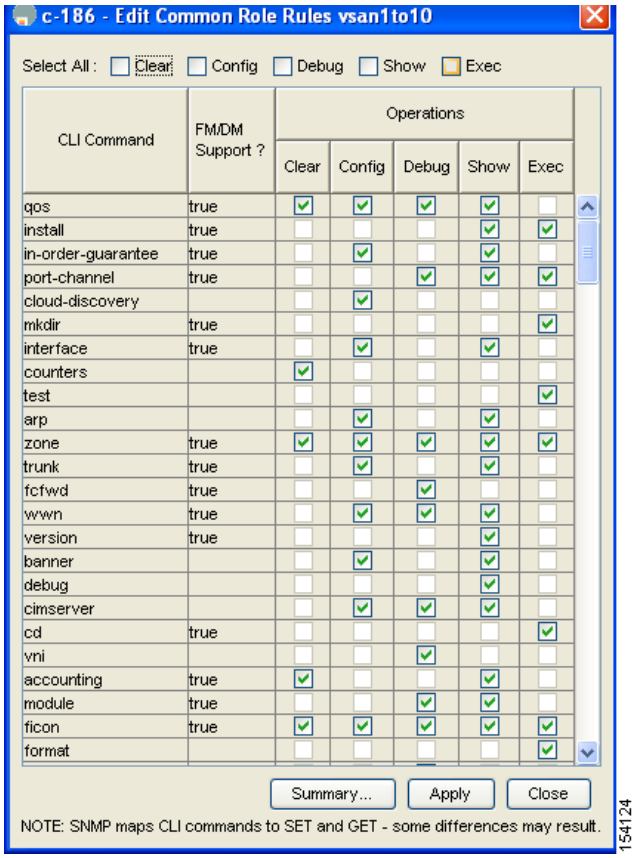

Step 5 Edit the rules you want to enable or disable for the common role.

Step 6 Click **Apply** to apply the new rules and close the Rules dialog, or click **Close** to close the Rules dialog without applying the rules.

Rule 1 is applied first, thus permitting, for example, sangroup users access to all **config** CLI commands. Rule 2 is applied next, denying FSPF configuration to sangroup users. As a result, sangroup users can perform all other **config** CLI commands, except **fspf** CLI configuration commands.

Note The order of rule placement is important. If you had swapped these two rules and issued the **deny config feature fspf** rule first and issued the **permit config** rule next, you would be allowing all sangroup users to perform all configuration commands because the second rule globally overrode the first rule.

Displaying Role-Based Information

The rules are displayed by rule number and are based on each role. All roles are displayed if the role name is not specified.

To view rules for a role using Device Manager, follow these steps:

Step 1 Click **Security > Roles.**

You see the Roles dialog box.

Step 2 Select a role name and click **Rules**.

You see the Rules dialog box.

Step 3 Click **Summary** to get a summarized view of the rules configured for this role.

Role Distributions

Role-based configurations use the Cisco Fabric Services (CFS) infrastructure to enable efficient database management, and to provide a single point of configuration for the entire fabric (see Chapter 12, “Using the CFS Infrastructure”).

The following configurations are distributed:

- Role names and descriptions
- List of rules for the roles
- VSAN policy and the list of permitted VSANs

This section includes the following topics:

About Role Databases

Role-based configurations use two databases to accept and implement configurations.

- Configuration database—The running database currently enforced by the fabric.
- Pending database—Your subsequent configuration changes are stored in the pending database. If you modify the configuration, you need to commit or discard the pending database changes to the configuration database. The fabric remains locked during this period. Changes to the pending database are not reflected in the configuration database until you commit the changes.

Locking the Fabric

The first action that modifies the database creates the pending database and locks the feature in the entire fabric. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database along with the first change.

Committing the Changes

If you commit the changes made to the pending database, the configuration is committed to all the switches in the fabric. On a successful commit, the configuration change is applied throughout the fabric and the lock is released. The configuration database now contains the committed changes and the pending database is now cleared.

To commit role-based configuration changes using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** in the Physical Attributes pane. Click the **Roles CFS** tab in the Information pane.

You see the screen shown in Figure 33-5.

Figure 33-5 Roles CFS Tab

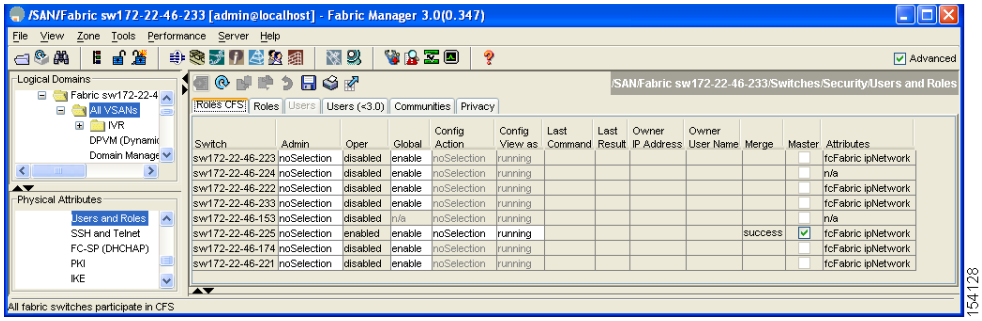

Step 2 Set the Global drop-down menu to **enable** to enable CFS.

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

Step 4 Set the Config Action drop-down menu to **commit** to commit the roles using CFS.

Step 5 Click **Apply Changes** to save this change.

Discarding the Changes

If you discard (abort) the changes made to the pending database, the configuration database remains unaffected and the lock is released.

To discard role-based configuration changes using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** in the Physical Attributes pane. Click the **Roles CFS** tab in the Information pane.

Step 2 Set the Config Action drop-down menu to **abort** to discard any uncommitted changes.

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

Enabling Distribution

To enable role-based configuration distribution using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** in the Physical Attributes pane. Click the **Roles CFS** tab in the Information pane.

Step 2 Set the Global drop-down menu to **enable** to enable CFS distribution.

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

Clearing Sessions

To forcibly clear the existing role session in the fabric using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** in the Physical Attributes pane. Click the **Roles CFS** tab in the Information pane.

Step 2 Set the Config Action drop-down menu to **clear** to clear the pending database.

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

Note Any changes in the pending database are lost when you clear a session.

Database Merge Guidelines

Fabric merge does not modify the role database on a switch. If two fabrics merge, and the fabrics have different role databases, the software generates an alert message.

See the "CFS Merge Support" section on page 12-9 for detailed concepts.

- Verify that the role database is identical on all switches in the entire fabric.
- Be sure to edit the role database on any switch to the desired database and then commit it. This synchronizes the role databases on all the switches in the fabric.

Displaying Roles When Distribution is Enabled

When you enable distribution for roles, you can view either the pending role database (the database before it is distributed) or the running database.

To view the roles using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** in the Physical Attributes pane. Click the **Roles CFS** tab in the Information pane (see Figure 33-6).

Figure 33-6 Roles CFS Tab

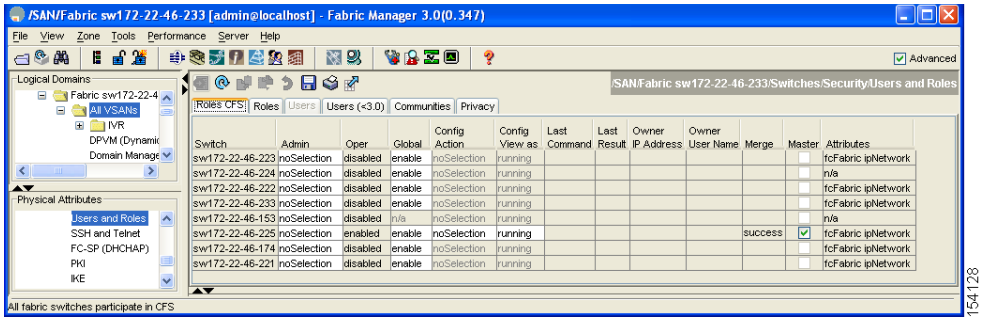

Step 2 Set the View Config As drop-down value to **pending** to view the pending database or set the View Config As drop-down menu to **running** to view the running database.

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

User Accounts

Every Cisco MDS 9000 Family switch user has the account information stored by the system. Your authentication information, user name, user password, password expiration date, and role membership are stored in your user profile.

The tasks explained in this section enable you to create users and modify the profile of an existing user. These tasks are restricted to privileged users as determined by your administrator.

The password should have the strong characteristics, such as the following:

- Are at least eight characters long
- Not contain many consecutive characters (such as “abcd”)
- Not contain many repeating characters (such as “aaabbb”)
- Not contain dictionary words
- Contain both upper- and lowercase characters
- Contain numbers

The following are examples of strong passwords:

- If2CoM18
- 2004AsdfLkj30
- Cb1955S21

Note Clear test passwords can only contain alphanumeric characters. Special characters such as the dollar sign ($) or the percent sign (%) are not allowed.

This section includes the following topics:

About Users

The passphrase specified in the **snmp-server user** option and the password specified **username** option are synchronized (see the “SNMPv3 CLI User Management and AAA Integration” section on page 34-3).

By default, the user account does not expire unless you explicitly configure it to expire. The **expire** option determines the date on which the user account is disabled. The date is specified in the YYYY-MM-DD format.

Tip The following words are reserved and cannot be used to configure users: bin, daemon, adm, lp, sync, shutdown, halt, mail, news, uucp, operator, games, gopher, ftp, nobody, nscd, mailnull, rpc, rpcuser, xfs, gdm, mtsuser, ftpuser, man, and sys.

Note User passwords are not displayed in the switch configuration file.

Tip If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password as shown in the sample configuration. Passwords are case-sensitive. “admin” is no longer the default password for any Cisco MDS 9000 Family switch. You must explicitly configure a strong password.

Caution Cisco MDS SAN-OS does not support all numeric user names, whether created with TACACS+ or RADIUS, or created locally. Local users with all numeric names cannot be created. If an all numeric user name exists on an AAA server and is entered during login, the user is not logged in.

Configuring Users

Note In Release 3.0(1) or later, you must be logged into Fabric Manager or Device Manager with your password and privacy password to create users.

To configure a new user or to modify the profile of an existing user using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane. Click the **Users** tab in the Information pane to see a list of users like the one in Figure 33-7.

Figure 33-7 Users listed under the Users Tab

Switch	User	Primary Role	Other Role	Digest	Encryption
sw172-22-46-182	admin	network-admin		MD5	NoPriv
sw172-22-46-224	abcd	network-operator	network-admin	MD5	DES
sw172-22-46-153	admin	network-admin		MD5	NoPriv
sw172-22-46-182	md5usr	network-operator		MD5	DES
sw172-22-46-224	admin	network-admin		MD5	DES
sw172-22-46-224	mchinn	network-operator	network-admin	MD5	DES
sw172-22-46-224	md5usr	network-admin		MD5	DES
sw172-22-46-153	mchinn	network-operator	network-admin	MD5	DES
sw172-22-46-224	mchinn1	network-operator	network-admin	MD5	NoPriv
sw172-22-46-224	shadmin	network-operator	network-admin	SHA	DES
sw172-22-46-153	md5usr	network-admin		MD5	DES
sw172-22-46-153	junknew	network-operator	network-admin	MD5	DES
sw172-22-46-153	mchinn1	network-operator	network-admin	MD5	DES
sw172-22-46-153	mchinn5	network-operator	network-admin	MD5	NoPriv
sw172-22-46-153	shadmin	network-operator	network-admin	SHA	DES
sw172-22-46-153	testUser	network-operator	test	MD5	DES

154126

Step 2 Click **Create Row**.

You see the Create Users dialog box shown in Figure 33-8.

Figure 33-8 Create Users Dialog Box

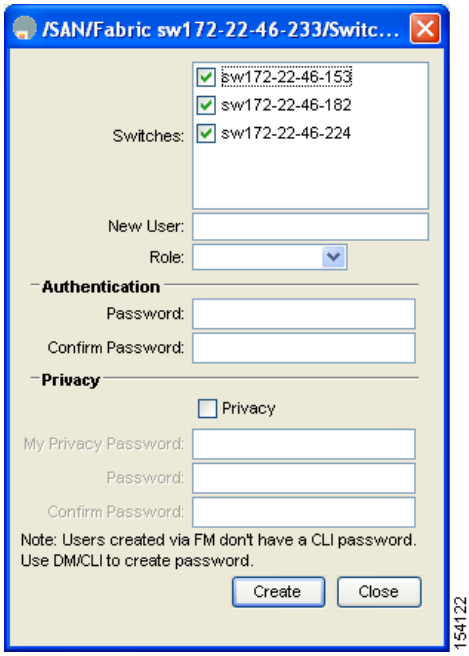

Step 3 Optionally alter the Switches check boxes to specify one or more switches.

Step 4 Enter the user name in the New User field.

Step 5 Select a role from the Role drop-down menu. You can also enter a new role name in the field if you do not want to select one from the drop-down menu. If you do this, go back and configure this role appropriately (see the "User Accounts" section on page 33-10).

Step 6 Enter the password for the user in the New Password and Confirm Password fields. Enter the same new password in the New Password and Confirm Password fields.

Step 7 Check the **Privacy** check box and complete the password fields to encrypt management traffic.

Step 8 Click **Create** to create the entry or click **Close** to discard any unsaved changes and close the dialog box.

Deleting a User

To delete a user using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane. Click the **Users** tab in the Information pane to see a list of users.

Step 2 Click the name of the user you want to delete.

Step 3 Click **Delete Row** to delete the selected user.

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

Displaying User Account Information

To display configured information about user accounts using Fabric Manager, follow these steps:

Step 1 Expand **Security** and then select **Users and Roles** in the Physical Attributes pane.

Step 2 Click the **Users** tab. You see the list of SNMP users shown in Figure 33-9 in the Information pane.

Figure 33-9 Users listed under the Users Tab

Switch	User	Primary Role	Other Role	Digest	Encryption
sw172-22-46-182	admin	network-admin		MD5	NoPriv
sw172-22-46-224	abcd	network-operator	network-admin	MD5	DES
sw172-22-46-153	admin	network-admin		MD5	NoPriv
sw172-22-46-182	md5usr	network-operator		MD5	DES
sw172-22-46-224	admin	network-admin		MD5	DES
sw172-22-46-224	mchinn	network-operator	network-admin	MD5	DES
sw172-22-46-224	md5usr	network-admin		MD5	DES
sw172-22-46-153	mchinn	network-operator	network-admin	MD5	DES
sw172-22-46-224	mchinn1	network-operator	network-admin	MD5	NoPriv
sw172-22-46-224	shadmin	network-operator	network-admin	SHA	DES
sw172-22-46-153	md5usr	network-admin		MD5	DES
sw172-22-46-153	junknew	network-operator	network-admin	MD5	DES
sw172-22-46-153	mchinn1	network-operator	network-admin	MD5	DES
sw172-22-46-153	mchinn5	network-operator	network-admin	MD5	NoPriv
sw172-22-46-153	shadmin	network-operator	network-admin	SHA	DES
sw172-22-46-153	testUser	network-operator	test	MD5	DES

154126

SSH Services

The Telnet service is enabled by default on all Cisco MDS 9000 Family switches. Before enabling the SSH service, generate a server key pair (see the "Generating the SSH Server Key Pair" section on page 33-16).

This section includes the following topics:

About SSH

SSH provides secure communications to the Cisco SAN-OS CLI. You can use SSH keys for the following SSH options:

- SSH1
- SSH2, using RSA
- SSH2 using DSA

About the SSH Server Key Pair

Be sure to have an SSH server key pair with the appropriate version before enabling the SSH service. Generate the SSH server key pair according to the SSH client version used. The number of bits specified for each key pair ranges from 768 to 2048.

The SSH service accepts three types of key pairs for use by SSH versions 1 and 2.

- The **rsa1** option generates the RSA1 key pair for the SSH version 1 protocol.
- The **dsa** option generates the DSA key pair for the SSH version 2 protocol.
- The **rsa** option generates the RSA key pair for the SSH version 2 protocol.

Caution If you delete all of the SSH keys, you cannot start a new SSH session.

Generating the SSH Server Key Pair

To generate the SSH server key pair, follow these steps:

Step 1 Expand **Switches > Security** and then select **SSH and Telnet**.

You see the configuration shown in Figure 33-10 in the Information pane.

Figure 33-10 SSH and Telnet Configuration

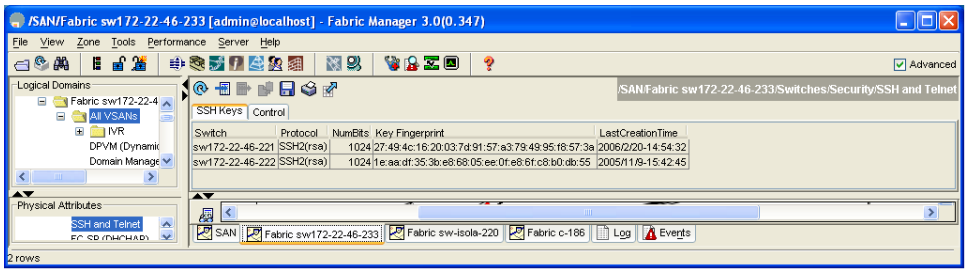

Step 2 Click **Create Row**.

You see the SSH and Telnet Key Create dialog box shown in Figure 33-11.

Figure 33-11 Create SSH and Telnet Dialog Box

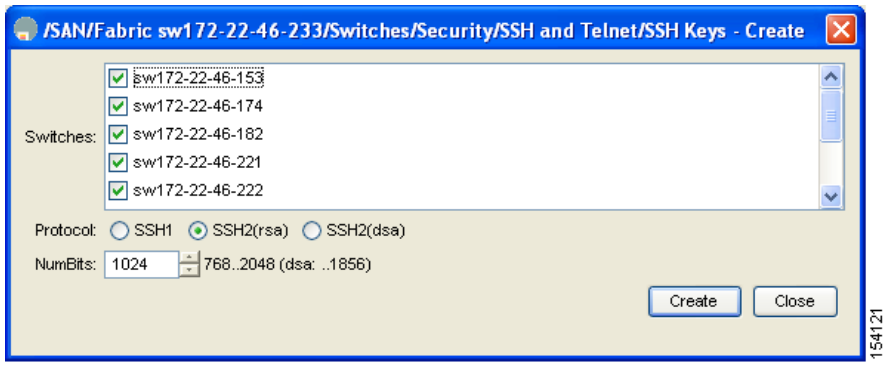

Step 3 Check the switches you want to assign to this SSH key pair.

Step 4 Choose the key pair option type from the listed Protocols. The listed protocols are SSH1, SSH2(rsa), and SSH2(dsa).

Step 5 Set the number of bits that will be used to generate the key pairs in the NumBits drop-down menu.

Step 6 Click **Create** to generate these keys or click **Close** to discard any unsaved changes.

Overwriting a Generated Key Pair

If the SSH key pair option is already generated for the required version, you can force the switch to overwrite the previously generated key pair.

To overwrite a previously generated key pair using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **SSH and Telnet**.

You see the configuration in the Information pane.

Step 2 Highlight the key that you want to overwrite and click **Delete Row**.

Step 3 Click **Apply Changes** to save these changes or click the **Undo Changes** to discard unsaved changes.

Step 4 Click the **Create Row**.

You see the SSH and Telnet Key Create dialog box.

Step 5 Check the switches you want to assign this SSH key pair.

Step 6 Choose the key pair option type from the Protocols radio buttons.

Step 7 Set the number of bits that will be used to generate the key pairs in the NumBits drop-down menu.

Step 8 Click **Create** to generate these keys or click **Close** to discard any unsaved changes.

Enabling SSH or Telnet Service

By default, the SSH service is disabled. Fabric Manager enables SSH automatically when you configure it.

To enable or disable SSH using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **SSH and Telnet**.

Step 2 Select the **Control** tab and check an **SSH** check box or **Telnet** check box for each switch as shown in Figure 33-12.

Figure 33-12 Control Tab under SSH and Telnet

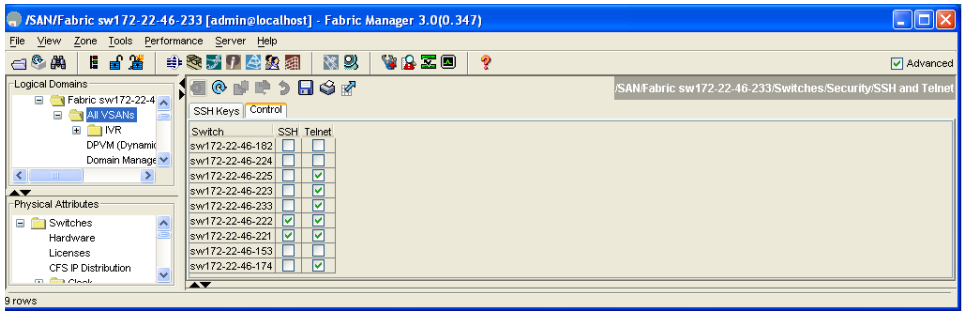

Step 3 Click **Apply Changes** to save this change or click **Undo Changes** to discard unsaved changes.

Note If you are logging in to a switch through SSH and you have issued the **aaa authentication login default none** CLI command, you must enter one or more key strokes to log in. If you press the **Enter** key without entering at least one keystroke, your log in will be rejected.

SSH Authentication Using Digital Certificates

SSH authentication on the Cisco MDS 9000 Family switches provide X.509 digital certificate support for host authentication. An X.509 digital certificate is a data item that vouches for the origin and integrity of a message. It contains encryption keys for secured communications and is "signed" by a trusted certification authority (CA) to verify the identity of the presenter. The X.509 digital certificate support provides either DSA or RSA algorithms for authentication.

The certificate infrastructure uses the first certificate that supports the Secure Socket Layer (SSL) and is returned by the security infrastructure, either through query or notification. Verification of certificates is successful if the certificates are from any of the trusted CAs.

You can configure your switch for either SSH authentication using an X.509 certificate or SSH authentication using a Public Key Certificate, but not both. If either of them is configured and the authentication fails, you will be prompted for a password.

For more information on CAs and digital certificates, see Chapter 38, "Configuring Certificate Authorities and Digital Certificates."

Creating or Updating Users

The passphrase specified in the **snmp-server user** option and the password specified **username** option are synchronized.

By default, the user account does not expire unless you explicitly configure it to expire. The **expire** option determines the date on which the user account is disabled. The date is specified in the YYYY-MM-DD format.

Tip The following words are reserved and cannot be used to configure users: bin, daemon, adm, lp, sync, shutdown, halt, mail, news, uucp, operator, games, gopher, ftp, nobody, nscd, mailnull, rpc, rpcuser, xfs, gdm, mtsuser, ftpuser, man, and sys.

Note User passwords are not displayed in the switch configuration file.

Tip If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password as shown in the sample configuration. Passwords are case-sensitive. "admin" is no longer the default password for any Cisco MDS 9000 Family switch. You must explicitly configure a strong password.

Caution Cisco MDS SAN-OS does not support all numeric usernames, whether created with TACACS+ or RADIUS, or created locally. Local users with all numeric names cannot be created. If an all numeric username exists on an AAA server and is entered during login, the user is not logged in.

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

Note Only the network-admin users are allowed to modify other user's privileges.

To configure a new user or to modify the profile of an existing user using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane. Click the **Users** tab in the Information pane to see the user information.

Step 2 Click **Create Row** to create a user.

You see the Create Users dialog box.

Step 3 Select the switches to which this user will be allowed access.

Step 4 Assign a new user name and password.

Note User account names must contain non-numeric characters.

Step 5 Select the roles that you want to assign to this new user.

Step 6 Select the digest and encryption for the user that you are creating or updating.

Step 7 Optionally, enter an expiry date and an SSH file name for the user.

Step 8 Click **Create** to create the user or **Close** to discard the changes.

Recovering the Administrator Password

You can recover the administrator password using one of two methods:

- From the CLI with a user name that has network-admin privileges.
- Power cycling the switch.

Note To recover an administrator's password, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

Configuring Cisco ACS Servers

The Cisco Access Control Server (ACS) uses TACACS+ and RADIUS protocols to provide AAA services that ensure a secure environment.When using the AAA server, user management is normally done using Cisco ACS. Figure 33-13, Figure 33-14, Figure 33-15, and Figure 33-16 display ACS server user setup configurations for network-admin roles and multiple roles using either TACACS+ or RADIUS.

Caution Cisco MDS SAN-OS does not support all numeric usernames, whether created with TACACS+ or RADIUS, or created locally. Local users with all numeric names cannot be created. If an all numeric user name exists on an AAA server and is entered during login, the user is not logged in.

Note Each role specified in the cisco-av-pair must exist in the MDS, or the user will have the 'network-operator' role.

Figure 33-13 Configuring the Network-admin Role When Using RADIUS

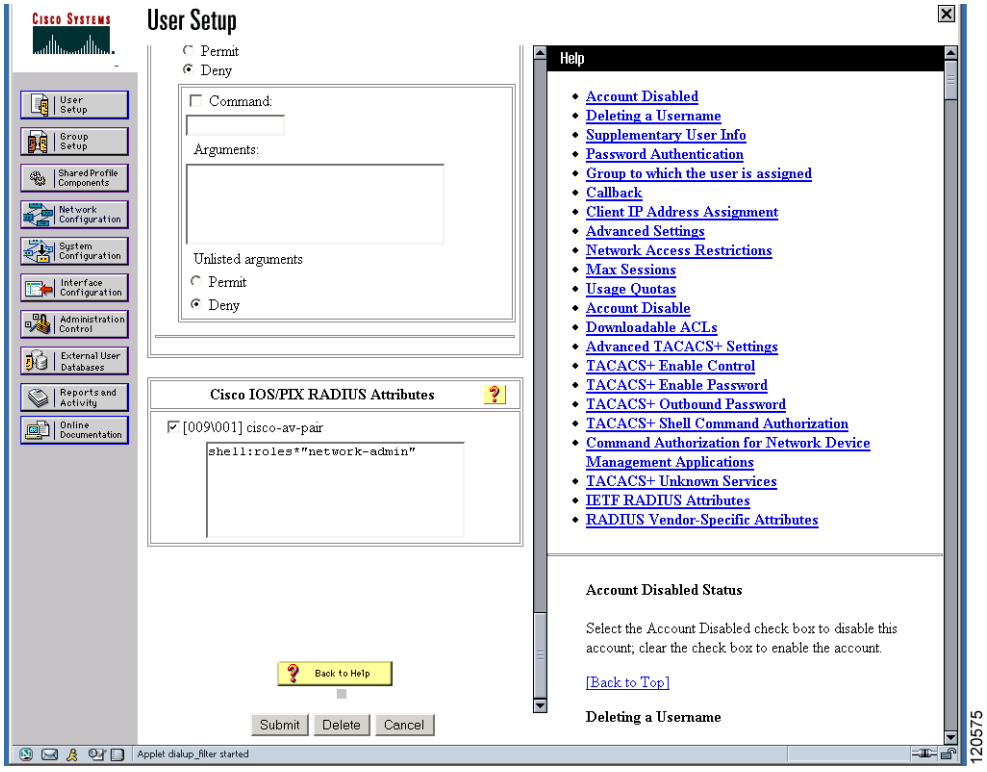

Figure 33-14 Configuring Multiple Roles with SNMPv3 Attributes When Using RADIUS

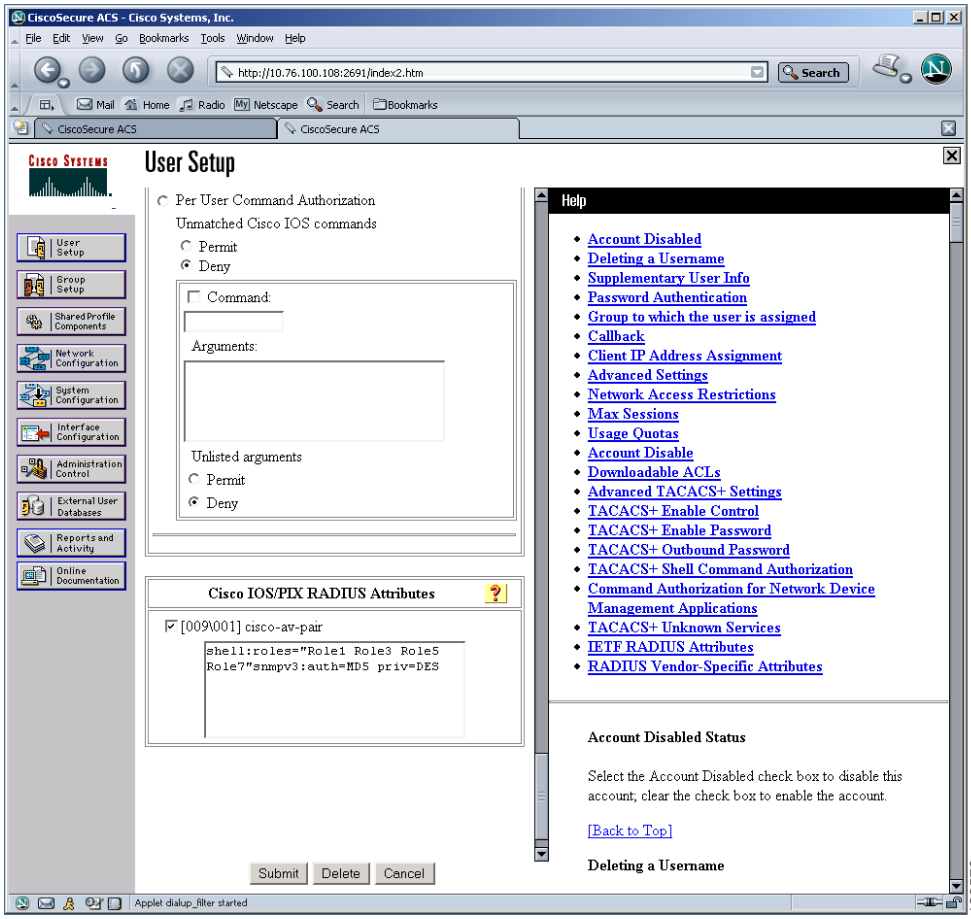

Figure 33-15 Configuring the network-admin Role with SNMPv3 Attributes When Using TACACS+

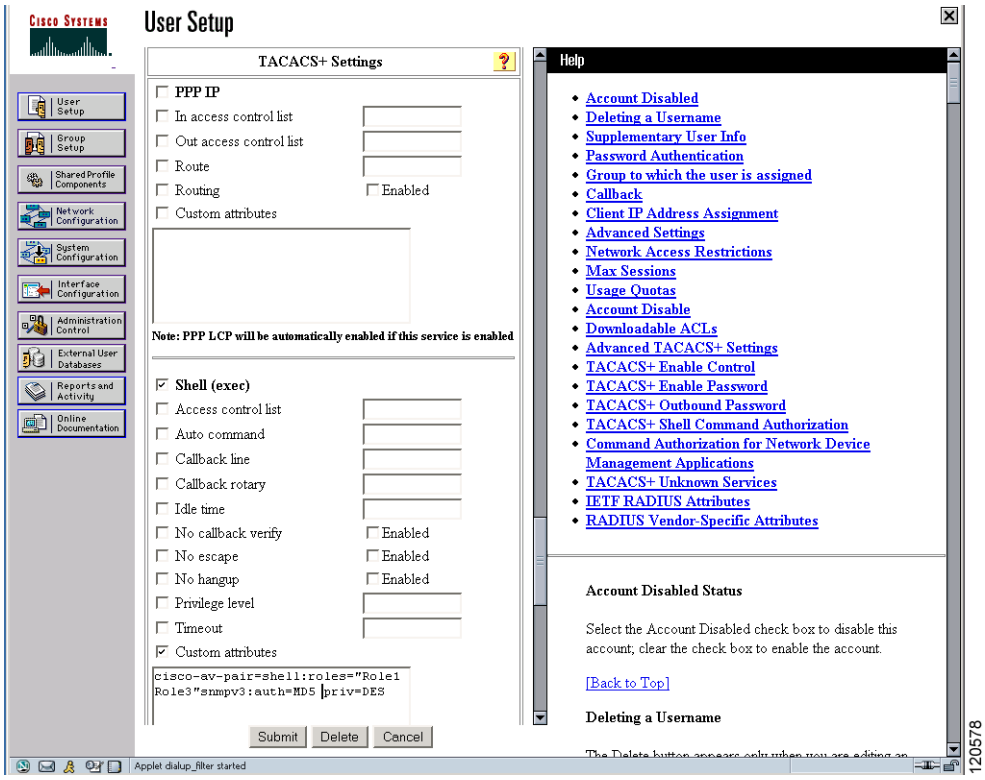

Figure 33-16 Configuring Multiple Roles with SNMPv3 Attributes When Using TACACS+

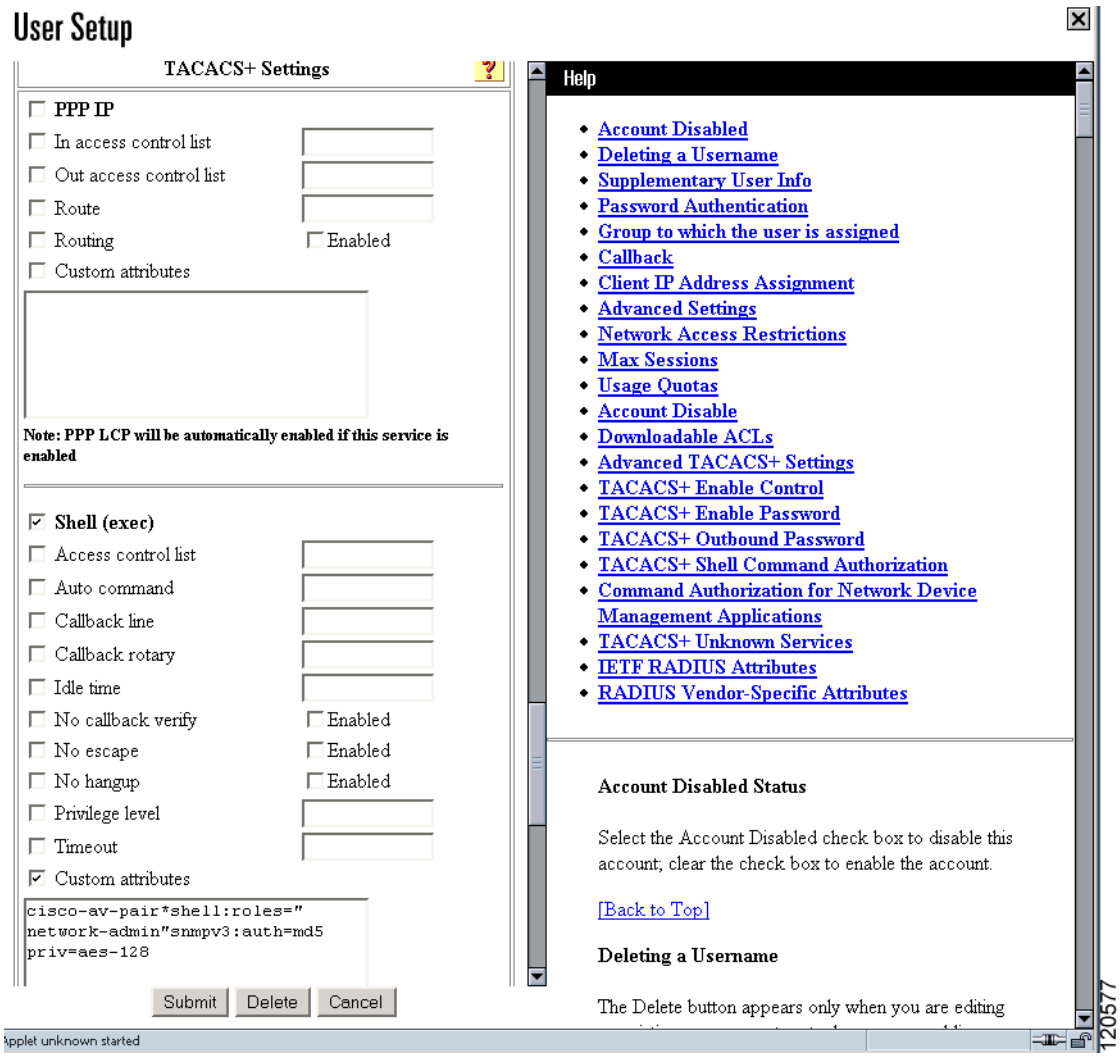

Default Settings

Table 33-1 lists the default settings for all switch security features in any switch.

Table 33-1 Default Switch Security Settings

Parameters	Default
Roles in Cisco MDS Switches	Network operator (network-operator).
AAA configuration services	Local.
Authentication port	1821.
Accounting port	1813.
Preshared key communication	Clear text.
RADIUS server time out	1 (one) second.

Table 33-1 Default Switch Security Settings (continued)

Parameters	Default
RADIUS server retries	Once.
TACACS+	Disabled.
TACACS+ servers	None configured.
TACACS+ server timeout	5 seconds.
AAA server distribution	Disabled.
VSAN policy for roles	Permit.
User account	No expiry (unless configured).
Password	None.
Accounting log size	250 KB.
SSH service	Disabled.
Telnet service	Enabled.

CHAPTER 34

Configuring SNMP

The CLI and SNMP use common roles in all switches in the Cisco MDS 9000 Family. You can use SNMP to modify a role that was created using the CLI and vice versa.

Users, passwords, and roles for all CLI and SNMP users are the same. A user configured through the CLI can access the switch using SNMP (for example, the Fabric Manager or the Device Manager) and vice versa.

This chapter includes the following sections:

About SNMP

SNMP is an application layer protocol that facilitates the exchange of management information between network devices. In all Cisco MDS 9000 Family switches, three SNMP versions are available: SNMPv1, SNMPv2c, and SNMPv3 (see Figure 34-1).

Figure 34-1 SNMP Security

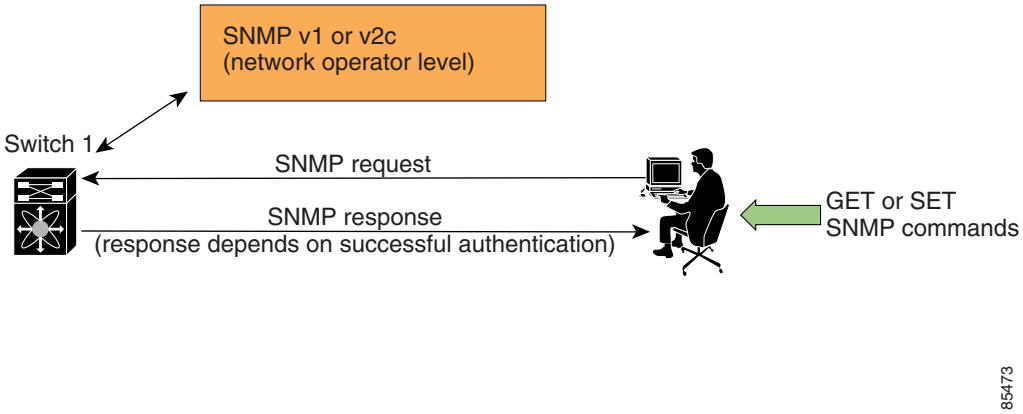

This section includes the following topics:

SNMP Version 1 and Version 2c

SNMP Version 1 (SNMPv1) and SNMP Version 2c (SNMPv2c) use a community string match for user authentication. Community strings provided a weak form of access control in earlier versions of SNMP. SNMPv3 provides much improved access control using strong authentication and should be preferred over SNMPv1 and SNMPv2c wherever it is supported.

SNMP Version 3

SNMP Version 3 (SNMPv3) is an interoperable standards-based protocol for network management. SNMPv3 provides secure access to devices by a combination of authenticating and encrypting frames over the network. The security features provided in SNMPv3 are:

- Message integrity—Ensures that a packet has not been tampered with in-transit.
- Authentication—Determines the message is from a valid source.
- Encryption—Scrambles the packet contents to prevent it from being seen by unauthorized sources.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

Assigning SNMP Switch Contact and Location Information

You can assign the switch contact information, which is limited to 32 characters (without spaces) and the switch location.

To configure contact and location information, using Fabric Manager, follow these steps:

Step 1 Expand **Switches** from the Physical Attributes pane. You see the switch settings in the Information pane.

Step 2 Fill in the Location and Contact fields for each switch.

Step 3 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

SNMPv3 CLI User Management and AAA Integration

The Cisco SAN-OS software implements RFC 3414 and RFC 3415, including user-based security model (USM) and role-based access control. While SNMP and the CLI have common role management and share the same credentials and access privileges, the local user database was not synchronized in earlier releases.

SNMPv3 user management can be centralized at the AAA server level. This centralized user management allows the SNMP agent running on the Cisco MDS switch to leverage the user authentication service of the AAA server. Once user authentication is verified, the SNMP PDUs are processed further. Additionally, the AAA server is also used to store user group names. SNMP uses the group names to apply the access/role policy that is locally available in the switch.

This section includes the following topics:

CLI and SNMP User Synchronization

Any configuration changes made to the user group, role, or password results in database synchronization for both SNMP and AAA.

Users are synchronized as follows:

- Deleting a user using results in the user being deleted for both SNMP and CLI.
- User-role mapping changes are synchronized in SNMP and CLI.

Note When the passphrase/password is specified in localized key/encrypted format, the password is not synchronized.

Note Starting in 3.0(1), the temporary SNMP login created for FM is no longer 24 hours. It is one hour.

- Existing SNMP users continue to retain the `auth` and `priv` passphrases without any changes.
- If the management station creates an SNMP user in the usmUserTable, the corresponding CLI user is created without any password (login is disabled) and will have the network-operator role.

Restricting Switch Access

You can restrict access to a Cisco MDS 9000 Family switch using IP Access Control Lists (IP-ACLs). See the "Configuring IPv4 Access Control Lists" section on page 36-1.

Group-Based SNMP Access

Note Because *group* is a standard SNMP term used industry-wide, we refer to role(s) as group(s) in this SNMP section.

SNMP access rights are organized by groups. Each group in SNMP is similar to a role through the CLI. Each group is defined with three accesses: read access, write access, and notification access. Each access can be enabled or disabled within each group.

You can begin communicating with the agent once your user name is created, your roles are set up by your administrator, and you are added to the roles.

Common Roles

CLI and SNMP in all switches in the Cisco MDS 9000 Family use common roles. You can use SNMP to modify a role that was created using CLI and vice versa (see Figure 34-2).

Figure 34-2 Common Roles

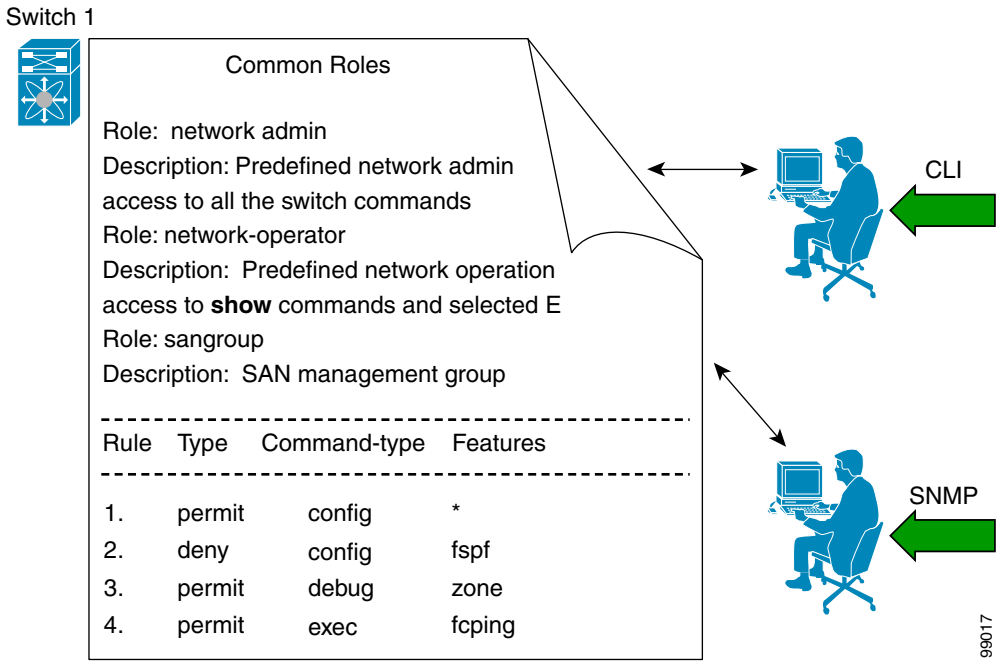

Each role in SNMP is the same as a role created or modified through the CLI (see the "Role-Based Authorization" section on page 33-1).

Each role can be restricted to one or more VSAN as required.

You can create new roles or modify existing roles using SNMP or the CLI.

- SNMP—Use the CISCO-COMMON-ROLES-MIB to configure or modify roles. Refer to the *Cisco MDS 9000 Family MIB Quick Reference*.
- CLI—Use the **role name** command.
- Fabric Manager—See the "Configuring Roles and Profiles" section on page 33-2.

SNMP has only three possible operations: GET, SET and NOTIFY. The CLI has five possible operations: DEBUG, SHOW, CONFIG, CLEAR and EXEC.

Note NOTIFY does not have any restrictions like the syslog messages in the CLI.

Table 34-1 explains how the CLI operations are mapped to the SNMP operations.

Table 34-1 CLI Operation to SNMP Operation Mapping

CLI Operation	SNMP Operation
DEBUG	Ignored
SHOW	GET
CONFIG	SET
CLEAR	SET
EXEC	SET

SNMP Users and Community Strings

You can create users or modify existing users using SNMP, Fabric Manager, or the CLI.

- SNMP—Create a user as a clone of an existing user in the usmUserTable on the switch. Once you have created the user, change the cloned secret key before activating the user. Refer to RFC 2574.
- Fabric Manager—See the "Configuring Users" section on page 33-12.
- CLI—Create a user or modify an existing user using the **snmp-server user** command.

By default only two roles are available in a Cisco MDS 9000 Family switch—network-operator and network-admin. You can also use any role that is configured in the Common Roles database (see the "Common Roles" section on page 34-5).

Tip All updates to the CLI security database and the SNMP user database are synchronized. You can use the SNMP password to log into either Fabric Manager or Device Manager. However, after you use the CLI password to log into Fabric Manager or Device Manager, you must use the CLI password for all future logins. If a user exists in both the SNMP database and the CLI database before upgrading to Cisco MDS SAN-OS Release 2.0(1b), then the set of roles assigned to the user becomes the union of both sets of roles after the upgrade.

This section includes the following topics:

About AES Encryption-Based Privacy

The Advanced Encryption Standard (AES) is the symmetric cipher algorithm. The Cisco SAN-OS software uses AES as one of the privacy protocols for SNMP message encryption and conforms with RFC 3826.

The **priv** option offers a choice of DES or 128-bit AES encryption for SNMP security encryption. The **priv** option along with the **aes-128** token indicates that this privacy password is for generating a 128-bit AES key.The AES priv password can have a minimum of eight characters. If the passphrases are specified in clear text, you can specify a maximum of 64 characters. If you use the localized key, you can specify a maximum of 130 characters.

Note For an SNMPv3 operation using the external AAA server, user configurations in the external AAA server require AES to be the privacy protocol to use SNMP PDU encryption.

Enforcing SNMPv3 Message Encryption

By default the SNMP agent allows the security parameters of authNoPriv and authPriv for the SNMPv3 messages that use user-configured SNMPv3 with auth and priv keys.

To enforce the message encryption for a user using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane Click the **Users** tab in the Information pane to see a list of users like the one shown in Figure 34-3.

Figure 34-3 User Information under the User Tab

Switch	User	Primary Role	Other Role	Digest	Encryption
sw172-22-46-182	admin	network-admin		MD5	NoPriv
sw172-22-46-224	abcd	network-operator	network-admin	MD5	DES
sw172-22-46-153	admin	network-admin		MD5	NoPriv
sw172-22-46-182	md5usr	network-operator		MD5	DES
sw172-22-46-224	admin	network-admin		MD5	DES
sw172-22-46-224	mchinn	network-operator	network-admin	MD5	DES
sw172-22-46-224	md5usr	network-admin		MD5	DES
sw172-22-46-153	mchinn	network-operator	network-admin	MD5	DES
sw172-22-46-224	mchinn1	network-operator	network-admin	MD5	NoPriv
sw172-22-46-224	shadmin	network-operator	network-admin	SHA	DES
sw172-22-46-153	md5usr	network-admin		MD5	DES
sw172-22-46-153	junknew	network-operator	network-admin	MD5	DES
sw172-22-46-153	mchinn1	network-operator	network-admin	MD5	DES
sw172-22-46-153	mchinn5	network-operator	network-admin	MD5	NoPriv
sw172-22-46-153	shadmin	network-operator	network-admin	SHA	DES
sw172-22-46-153	testUser	network-operator	test	MD5	DES

154126

Step 2 Click **Create Row**.

You see the Create Users dialog box.

The dialog box from Fabric Manager also provides check boxes to specify one or more switches.

Step 3 Enter the user name in the **New User** field.

Step 4 Select the role from the Role drop-down menu. You can enter a new role name in the field if you do not want to select one from the drop-down menu. If you do this, you must go back and configure this role appropriately (see the "User Accounts" section on page 33-10).

Step 5 Enter the same password for the user in both the New Password and Confirm Password fields.

Step 6 Check the **Privacy** check box and complete the password fields to encrypt management traffic.

Step 7 Click **Create** to create the new entry or click **Close** to discard any unsaved changes and close the dialog box.

To enforce the SNMPv3 message encryption globally on all the users using Fabric Manager, follow these steps:

Step 1 Select a VSAN in the Logical Domains pane. This will not work if you select All VSANS.

Step 2 Expand **Switches > Security** and then select **Users and Roles** in the Physical Attributes pane. Click the **Global** tab in the Information pane.

Step 3 Check the **GlobalEnforcePriv** check box.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard unsaved changes.

Assigning SNMPv3 Users to Multiple Roles

The SNMP server user configuration is enhanced to accommodate multiple roles (groups) for SNMPv3 users. After the initial SNMPv3 user creation, you can map additional roles for the user.

Note Only users belonging to a network-admin role can assign roles to other users.

To add multiple roles to a new user using Device Manager, follow these steps:

Step 1 Click **Security > Users** and the click the **Users** tab

You see the information shown in Figure 34-4.

Figure 34-4 Users Information in Device Manager

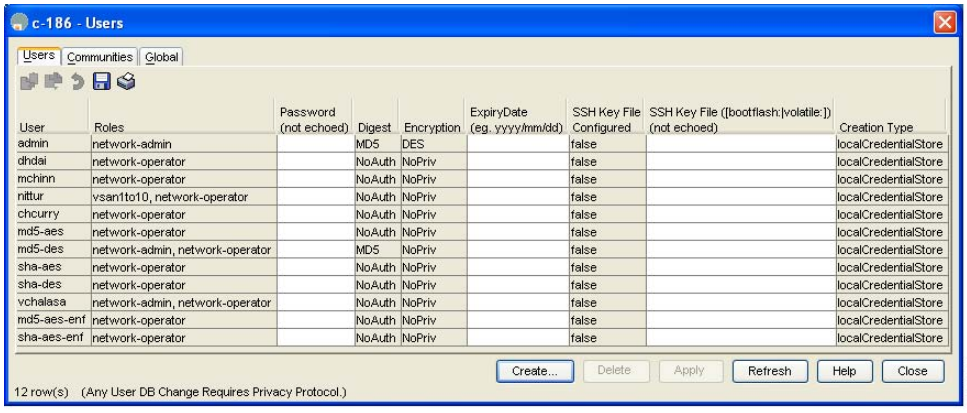

Step 2 Click **Create**.

Step 3 Enter the user name and password.

Step 4 Choose roles using the check boxes.

Step 5 Choose an option for Digest and one for Encryption.

Step 6 Optionally provide an expiration date for the user and the file name of an SSH key.

Step 7 Click **Create** to create the new entry or click **Close** close the dialog box without creating an entry.

Creating an SNMP Community String

You can configure read-only or read-write access for SNMPv1 and SNMPv2 users. Refer to RFC 2576.

To create an SNMPv1 or SNMPv2c community string using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane. Click the **Communities** tab in the Information pane.

You see existing communities (see Figure 34-5).

Figure 34-5 Communities Tab under Users and Roles

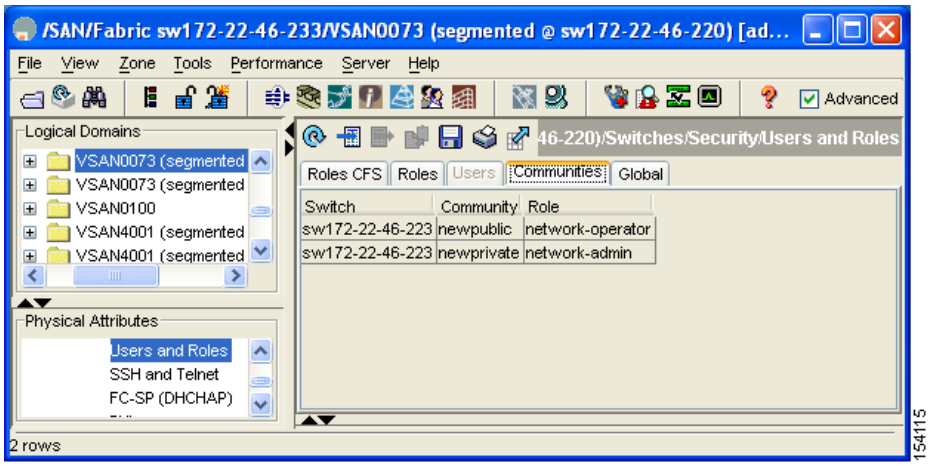

Step 2 Click **Create Row**.

You see the Create Community String dialog box. See Figure 34-6.

Figure 34-6 Create Community String Dialog Box

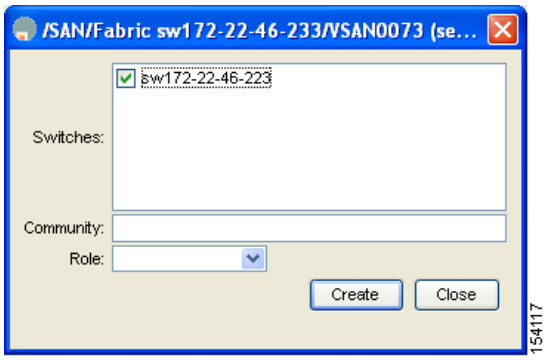

Step 3 Check the **Switch** check boxes to specify one or more switches.

Step 4 Enter the community name in the Community field.

Step 5 Select the role from Role drop-down list .

Note You can enter a new role name in the field if you do not want to select one from the drop-down list. If you do this, you must go back and configure this role appropriately (see the "Role-Based Authorization" section on page 33-1).

Step 6 Click **Create** to create the new entry or click **Close** to create the entry and close the dialog box.

Deleting a Community String

To delete a community string using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane. Click the **Communities** tab in the Information pane.

Step 2 Click the name of the community you want to delete.

Step 3 Click **Delete Row** to delete this community.

Viewing SNMP Community and User Information

To view information about SNMP users, roles, and communities from Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** in the Physical Attributes pane.

Step 2 Click either the **Users, Roles,** or **Communities** tab.

You see the list of SNMP users, roles, or communities in the Information pane. The communities are shown in Figure 34-7.

Figure 34-7 Communities Tab under Users and Roles

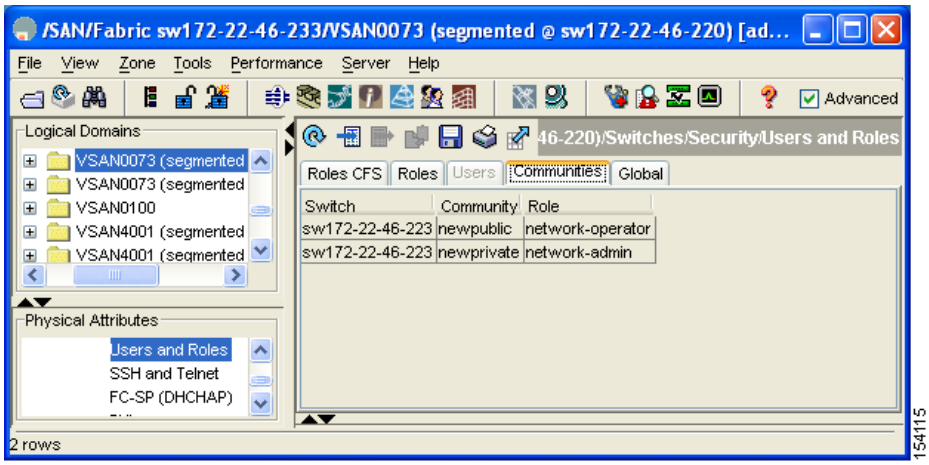

SNMP Trap and Inform Notifications

You can configure the Cisco MDS switch to send notifications to SNMP managers when particular events occur. You can send these notifications as traps or informs. Traps are unreliable because the receiver does not send any acknowledgment when it receives a trap. The sender cannot determine if the trap was received. However, an SNMP manager that receives an inform acknowledges the message with an SNMP response PDU. If the sender never receives a response, the inform is normally retransmitted. Thus, informs are more likely to reach their intended destinations.

Note Use the SNMP-TARGET-MIB to obtain more information on the destinations to which notifications are to be sent either as traps or as informs. Refer to the *Cisco MDS 9000 Family MIB Quick Reference*.

Tip SNMPv1 does not support informs.

This section includes the following topics:

- Configuring SNMPv1 and SNMPv2c Notifications, page 34-12
- Configuring SNMPv3 Notifications, page 34-13
- Enabling SNMP Notifications, page 34-13
- Configuring the Notification Target User for Informs, page 34-14
- Configuring SNMP Event Security, page 34-15
- Viewing the SNMP Events Log, page 34-15

Configuring SNMPv1 and SNMPv2c Notifications

To configure SNMPv1 and SNMPv2c notifications) using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Events** and then select **SNMP Traps** in the Physical Attributes pane.

You see the SNMP notification configuration in the Information pane shown in Figure 34-8.

Figure 34-8 SNMP Notifications

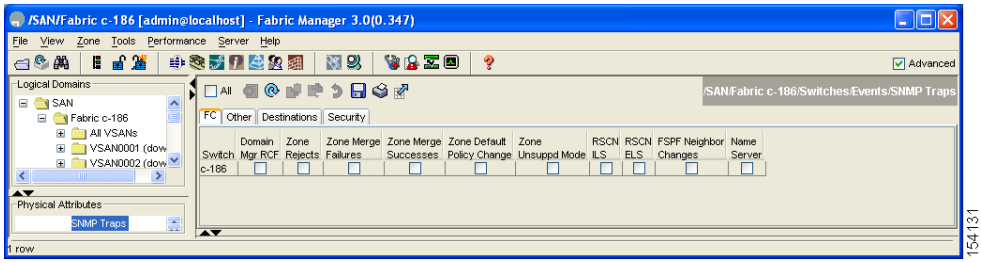

Step 2 Click the **Destinations** tab to add or modify a receiver for SNMP notifications.

Step 3 Click **Create Row** to create a new notification destination.

You see the Create Destination dialog box.

Step 4 Check the switches for which you want to configure a new destination.

Step 5 Set the destination IP address and UDP port.

Step 6 Set the Security to **v1** or **v2c**. (Note that **v3** security is used for SNMPv3, described separately.)

Step 7 Choose either the **trap** or **inform** radio button.

Step 8 Optionally, set the timeout or retry count values.

Step 9 Click **Create** to add this destination to the selected switches or click **Close** to discard any unsaved changes.

Step 10 Optionally, click the other tabs to enable specific notification types per switch.

Step 11 Click **Apply changes** to create the entry or click **Undo Changes** to discard any unsaved changes.

Configuring SNMPv3 Notifications

To configure SNMPv3 notifications (traps or informs) using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Events** and then select **SNMP Traps** in the Physical Attributes pane.

You see the SNMP configuration in the Information pane

Step 2 Click the **Destinations** tab to add or modify a receiver for SNMP notifications.

Step 3 Click **Create Row** to create a new notification destination.

You see the Create Destination dialog box

Step 4 Check the switches for which you want to configure a new destination.

Step 5 Set the destination IP address and UDP port.

Step 6 Set the Security to **v3**.

Step 7 Choose either the **trap** or **inform** radio button.

Step 8 Optionally, set the inform time out and retry values.

Step 9 Click **Create** to add this destination to the selected switches or click **Close** to discard any unsaved changes.

Step 10 Optionally, click the other tabs to enable specific notification types per switch.

Step 11 Click **Apply changes** to create the entry or click **Undo Changes** to discard any unsaved changes.

Note In the case of SNMPv3 notifications, the SNMP manager is expected to know the user credentials (authKey/PrivKey) based on the switch's engine ID to authenticate and decrypt the SNMP messages.

Enabling SNMP Notifications

Notifications (traps and informs) are system alerts that the switch generates when certain events occur. By default, no notification is defined or issued. If a notification name is not specified, all notifications are disabled or enabled.

Table 34-2 lists the Fabric Manager procedures that enable the notifications for Cisco MDS MIBs. Choose **Switches > Events > SNMP Traps** to see the check boxes listed in this table.

Note Choosing **Switches > Events > SNMP Traps** enables both traps and informs, depending on how you configured notifications. See the Configuring SNMPv1 and SNMPv2c Notifications, page 34-12 or the Configuring SNMPv3 Notifications, page 34-13.

Table 34-2 Enabling SNMP Notifications

MIB	Fabric Manager Check boxes
CISCO-ENTITY-FRU-CONTROL-MIB	Select the **Other** tab and check **FRU Changes.**
CISCO-FCC-MIB	Select the **Other** tab and check **FCC.**

Table 34-2 Enabling SNMP Notifications

MIB	Fabric Manager Check boxes
CISCO-DM-MIB	Select the **FC** tab and check **Domain Mgr RCF.**
CISCO-NS-MIB	Select the **FC** tab and check **Name Server.**
CISCO-FCS-MIB	Select the **Othe**r tab and check **FCS Rejects**
CISCO-FDMI-MIB	Select the **Othe**r tab and check **FDMI.**
CISCO-FSPF-MIB	Select the **FC** tab and check **FSPF Neighbor Change.**
CISCO-LICENSE-MGR-MIB	Select the **Othe**r tab and check **License Manager.**
CISCO-IPSEC-SIGNALLING-MIB	Select the **Othe**r tab and check **IPSEC.**
CISCO-PSM-MIB	Select the **Othe**r tab and check **Port Security.**
CISCO-RSCN-MIB	Select the **FC** tab and check **RSCN ILS,** and **RCSN ELS.**
SNMPv2-MIB	Select the **Othe**r tab and check **SNMP AuthFailure.**
VRRP-MIB, CISCO-IETF-VRRP-MIB	Select the **Othe**r tab and check **VRRP.**
CISCO-ZS-MIB	Select the **FC** tab and check **Zone Rejects, Zone Merge Failures, Zone Merge Successes, Zone Default Policy Change,** and **Zone Unsuppd Mode.**

The following notifications are enabled by default:

- **entity fru**
- **license**
- **link ietf-extended**

All other notifications are disabled by default.

To enable individual notifications using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Events** and then select **SNMP Traps** in the Physical Attributes pane.

You see the SNMP notification configuration in the Information pane.

Step 2 Click the **FC** tab to enable Fibre Channel related notifications.

Step 3 Check each notification check box that you want to enable.

Step 4 Click the **Other** tab to enable other notifications.

Step 5 Check each notification check box that you want to enable.

Step 6 Click **Apply changes** to create the entry or click **Undo Changes** to discard any unsaved changes.

Configuring the Notification Target User for Informs

You must configure a notification target user on the switch for sending SNMPv3 inform notifications to the SNMP manager.

To configure the notification target user, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

The credentials of the notification target user are used for encrypting the SNMPv3 inform notification messages to the configured SNMP manager.

Note For authenticating and decrypting the received INFORM PDU, the SNMP manager should have the same user credentials in its local configuration data store of users.

Configuring SNMP Event Security

Caution This is an advanced function that should only be used by administrators having experience with SNMPv3.

SNMP events can be secured against interception or eavesdropping in the same way that SNMP messages are secured. Fabric Manager or Device Manager allow you to configure the message processing model, the security model, and the security level for the SNMP events that the switch generates.

To configure SNMP event security using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Events** and then select **SNMP Traps. C**lick the **Security** tab in the Information pane.

You see the security information for SNMP notifications.

Step 2 Set the message protocol model (MPModel), security model, security name, and security level.

Step 3 Click **Apply Changes** to save and apply your changes.

Viewing the SNMP Events Log

To view the SNMP events log from Device Manager, click **Logs > FM Server > Events > Current** or **Logs > Events > Older** (see Figure 34-9). You see the Events Log dialog box with a log of events for a single switch.

Figure 34-9 Device Manager Events

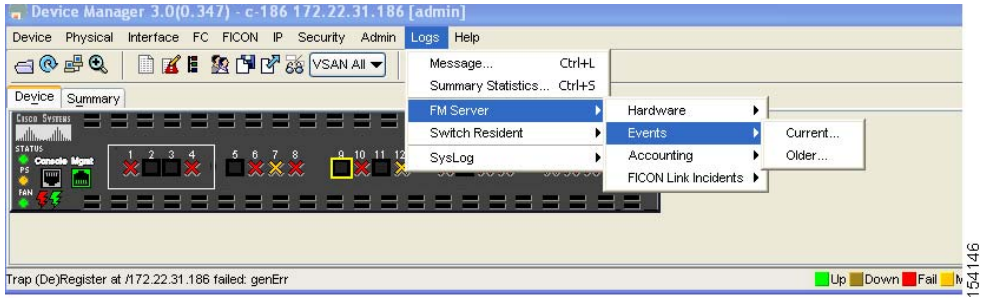

Note The MDS syslog manager must be set up before you can view the event logs.

Caution Changing these values from different Fabric Manager workstations at the same time may cause unpredictable results.

Default Settings

Table 34-3 lists the default settings for all SNMP features in any switch.

Table 34-3 Default SNMP Settings

Parameters	Default
User account	No expiry (unless configured).
Password	None.

CHAPTER 35

Configuring RADIUS and TACACS+

The authentication, authorization, and accounting (AAA) mechanism verifies the identity of, grants access to, and tracks the actions of users managing a switch. All Cisco MDS 9000 Family switches use Remote Access Dial-In User Service (RADIUS) and Terminal Access Controller Access Control System Plus (TACACS+) protocols to provide solutions using remote AAA servers.

Based on the user ID and password combination provided, switches perform local authentication or authorization using the local database or remote authentication or authorization using AAA server(s). A preshared secret key provides security for communication between the switch and AAA servers. This secret key can be configured for all AAA server or for only a specific AAA server. This security mechanism provides a central management capability for AAA servers.

This chapter includes the following sections:

Switch Management Security

Management security in any switch in the Cisco MDS 9000 Family provides security to all management access methods including the command-line interface (CLI) or Simple Network Management Protocol (SNMP).

This section includes the following topics:

Fabric Manager Security Options

You can access Fabric Manager using the console (serial connection), Telnet, or Secure Shell (SSH). For each management path (console or Telnet and SSH), you can configure one or more of the following security control options: local, remote (RADIUS or TACACS+), or none.

- Remote security control
  - Using Remote Authentication Dial-In User Services (RADIUS). See the "Configuring RADIUS Server Monitoring Parameters" section on page 35-8.
  - Using Terminal Access Controller Access Control System plus (TACACS+). See the "Configuring TACACS+ Server Monitoring Parameters" section on page 35-14.
- Local security control. See the "MSCHAP Authentication" section on page 35-26.

These security mechanisms can also be configured for the following scenarios:

- iSCSI authentication (see the "iSCSI Authentication Setup Guidelines and Scenarios" section on page 45-64).
- Fibre Channel Security Protocol (FC-SP) authentication (see the Chapter 40, "Configuring FC-SP and DHCHAP")

SNMP Security Options

The SNMP agent supports security features for SNMPv1,SNMPv 2c, and SNMPv3. Normal SNMP security mechanisms apply to all applications that use SNMP (for example, Cisco MDS 9000 Fabric Manager).

CLI security options also apply to the Fabric Manager and Device Manager.

See Chapter 34, "Configuring SNMP".

Switch AAA

Using the CLI or Fabric Manager, you can configure authentication, authorization, and accounting (AAA) switch functionalities on any switch in the Cisco MDS 9000 Family.

This section includes the following topics:

Authentication

Authentication is the process of verifying the identity of the person managing the switch. This identity verification is based on the user ID and password combination provided by the person trying to manage the switch. Cisco MDS 9000 Family switches allow you to perform local authentication (using the local lookup database) or remote authentication (using one or more RADIUS or TACACS+ servers).

Note When you log in to a Cisco MDS switch successfully using the Fabric Manager or Device Manager through Telnet or SSH and if that switch is configured for AAA server-based authentication, a temporary SNMP user entry is automatically created with an expiry time of one day. The switch traps the SNMPv3 protocol data units (PDUs) with your Telnet or SSH login name as the SNMPv3 user. The management station can temporarily use the Telnet/SSH login name as the SNMPv3 `auth` and `priv` passphrase. This temporary SNMP login is only allowed if you have one or more active MDS shell sessions. If you do not have an active session at any given time, your login is deleted and you will not be allowed to perform SNMPV3 operations.

Note Fabric Manager does not support AAA passwords with trailing white space, for example “passwordA “.

Authorization

By default, two roles exist in all Cisco MDS switches:

- Network operator (network-operator)—Has permission to view the configuration only. The operator cannot make any configuration changes.
- Network administrator (network-admin)— Has permission to execute all commands and make configuration changes. The administrator can also create and customize up to 64 additional roles.

If you use a SAN Volume Controller (SVC) setup, two more default roles exist in all Cisco MDS switches:

- SVC administrator (svc-admin)— Has permission to view the entire configuration and make SVC-specific configuration changes within the `switch(svc)` prompt.
- SVC operator (svc-operator)—Has permission to view the entire configuration. The operator cannot make any configuration changes.

Note Refer to the *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for more information on SVC.

These four default roles cannot be changed or deleted. You can create additional roles and configure the following options:

- Configure role-based authorization by assigning user roles locally or using remote AAA servers.
- Configure user profiles on a remote AAA server to contain role information. This role information is automatically downloaded and used when the user is authenticated through the remote AAA server.

Note If a user only belongs to one of the newly created roles and that role is subsequently deleted, then the user immediately defaults to the network-operator role.

Accounting

The accounting feature tracks and maintains a log of every management session used to access the switch. This information can be used to generate reports for troubleshooting and auditing purposes. Accounting logs can be stored locally or sent to remote AAA servers.

Remote AAA Services

Remote AAA services provided through RADIUS and TACACS+ protocols have the following advantages over AAA servers:

- It is easier to manage user password lists for each switch in the fabric.
- AAA servers are already deployed widely across enterprises and can be easily adopted.
- The accounting log for all switches in the fabric can be centrally managed.
- It is easier to manage user role mapping for each switch in the fabric.

Remote Authentication Guidelines

If you prefer using remote AAA servers, follow these guidelines:

- A minimum of one AAA server should be IP reachable.
- Be sure to configure a desired local AAA policy as this policy is used if all AAA servers are not reachable.
- AAA servers are easily reachable if an overlay Ethernet LAN is attached to the switch (see Chapter 47, “Configuring IP Storage”). We recommend this method.
- SAN networks connected to the switch should have at least one gateway switch connected to the Ethernet LAN reaching the AAA servers.

Server Groups

You can specify remote AAA servers for authentication, authorization, and accounting using server groups. A server group is a set of remote AAA servers implementing the same AAA protocol. The purpose of a server group is to provide for fail-over servers in case a remote AAA server fails to respond. If the first remote server in the group fails to respond, the next remote server in the group is tried until one of the servers sends a response. If all the AAA servers in the server group fail to respond, then that server group option is considered a failure. If required, you can specify multiple server groups. If the Cisco MDS switch encounters errors from the servers in the first group, it tries the servers in the next server group.

AAA Configuration Options

AAA configuration in Cisco 9000 Family switches is service based. You can have separate AAA configurations for the following services

- Telnet or SSH login (Fabric Manager and Device Manager login).
- Console login.
- iSCSI authentication (see the "iSCSI Authentication Setup Guidelines and Scenarios" section on page 45-64).
- FC-SP authentication (see Chapter 40, "Configuring FC-SP and DHCHAP").
- Accounting.

In general, server group, local, and none are the three options that can be specified for any service in an AAA configuration. Each option is tried in the order specified. If all the methods fail, local is tried.

Caution Cisco MDS SAN-OS does not support all numeric user names, whether created with TACACS+ or RADIUS, or created locally. Local users with all numeric names cannot be created. If an all numeric user name exists on an AAA server and is entered during login, the user is not logged in.

Note Even if local is not specified as one of the options, it is tried when all other configured options fail.

When RADIUS times out, local login is always attempted. For this local login to be successful, a local account for the user with the same password should exist and the RADIUS timeout and retries should take less than 40 seconds. The user is authenticated if the username and password exists in the local authentication configuration.

AAA Server Monitoring

An unresponsive AAA server introduces delay in processing of AAA requests. An MDS switch can periodically monitor an AAA server to check whether it is responding (or alive) to save time in processing AAA requests. The MDS switch marks unresponsive AAA servers as dead and does not send AAA requests to any dead AAA servers. An MDS switch periodically monitors dead AAA servers and brings them to the alive state once they are responding. This monitoring process verifies that an AAA server is in a working state before real AAA requests are sent its way. Whenever an AAA server changes to the dead or alive state, an SNMP trap is generated and the MDS switch warns the administrator that a failure is taking place before it can impact performance. See Figure 35-1 on page 35-6.

Figure 35-1 AAA Server States

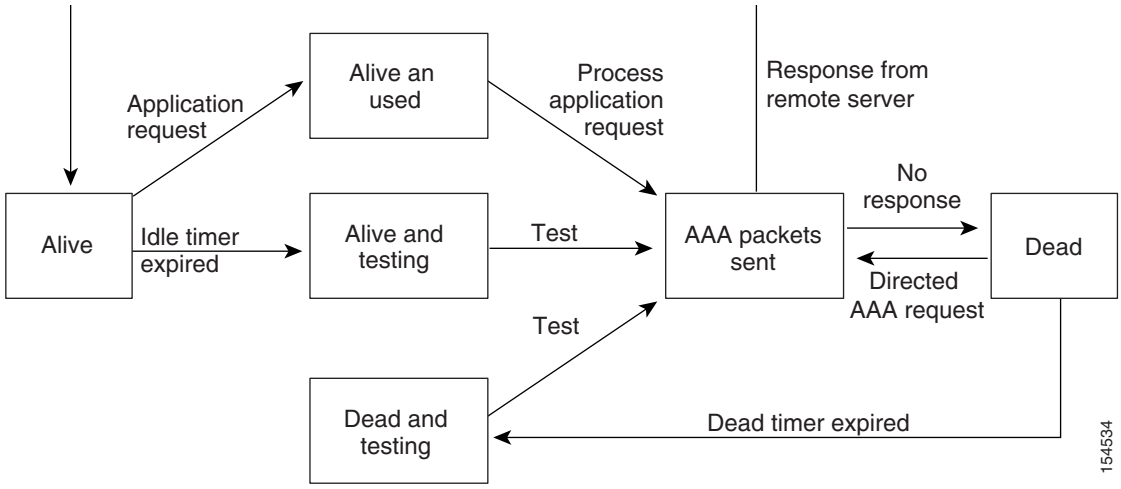

Note The monitoring interval for alive servers and dead servers are different and can be configured by the user. The AAA server monitoring is performed by sending a test authentication request to the AAA server.

The user name and password to be used in the test packet can be configured.

See the "Configuring RADIUS Server Monitoring Parameters" section on page 35-8 and "Configuring TACACS+ Server Monitoring Parameters" section on page 35-14.

Authentication and Authorization Process

Authentication is the process of verifying the identity of the person managing the switch. This identity verification is based on the user ID and password combination provided by the person managing the switch. The Cisco MDS 9000 Family switches allow you to perform local authentication (using the lookup database) or remote authentication (using one or more RADIUS servers or TACACS+ servers).

Figure 35-2 shows a flow chart of the process. The following steps explain the authorization and authentication process.

Step 1 When you can log in to the required switch in the Cisco MDS 9000 Family, you can use the Telnet, SSH, Fabric Manager/Device Manager, or console login options.

Step 2 When you have configured server groups using the server group authentication method, an authentication request is sent to the first AAA server in the group.

- If the AAA server fails to respond, then the next AAA server is tried and so on until the remote server responds to the authentication request.
- If all AAA servers in the server group fail to respond, then the servers in the next server group are tried.
- If all configured methods fail, then the local database is used for authentication.

Step 3 If you are successfully authenticated through a remote AAA server, then the following possibilities apply.

- If AAA server protocol is RADIUS, then user roles specified in the `cisco-av-pair` attribute are downloaded with an authentication response.
- If AAA server protocol is TACACS+, then another request is sent to the same server to get the user roles specified as custom attributes for the shell.
- If user roles are not successfully retrieved from the remote AAA server, then the user is assigned the network-operator role.

Step 4 If your user name and password are successfully authenticated locally, you are allowed to log in, and you are assigned the roles configured in the local database.

Figure 35-2 Switch Authorization and Authentication Flow

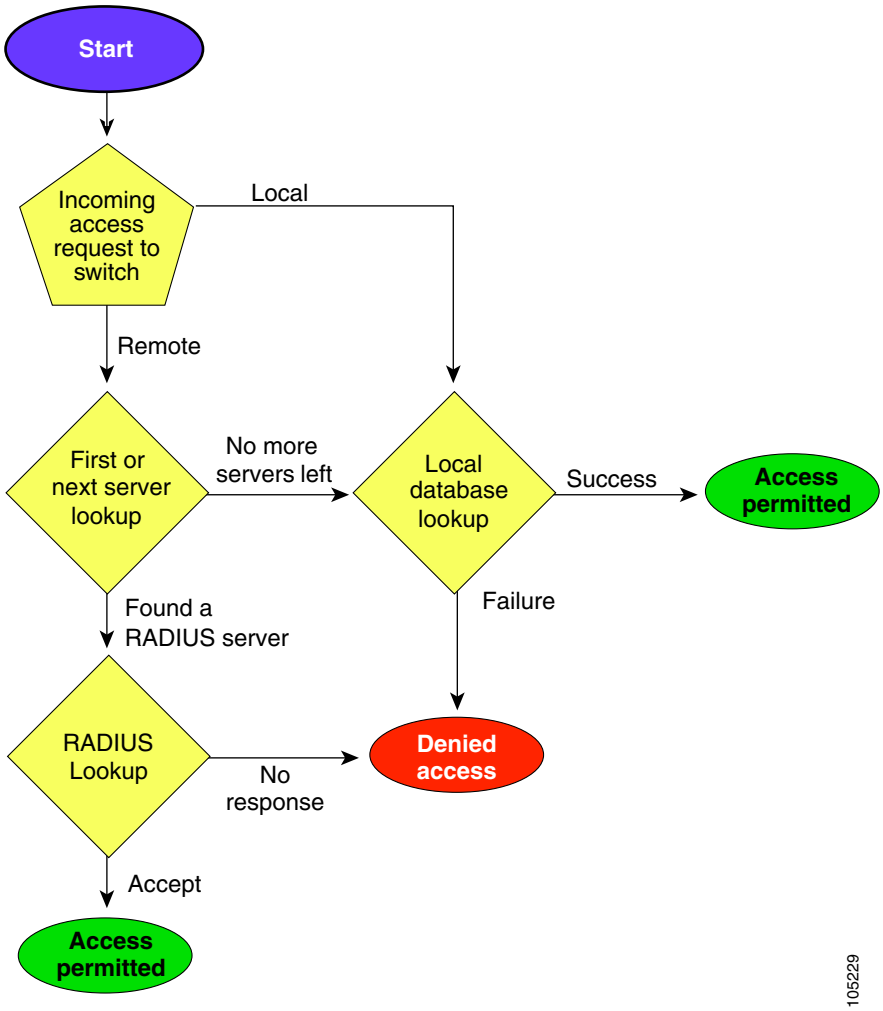

Note No more server groups left = no response from any server in all server groups.
No more servers left = no response from any server within this server group.

Configuring RADIUS Server Monitoring Parameters

Cisco MDS 9000 Family switches can use the RADIUS protocol to communicate with remote AAA servers. You can configure multiple RADIUS servers and server groups and set timeout and retry counts.

This section defines the RADIUS operation, identifies its network environments, and describes its configuration possibilities.

RADIUS is a distributed client/server protocol that secures networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco MDS 9000 Family switches and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.

This section includes the following topics:

About RADIUS Server Default Configuration

Fabric Manager allows you to set up a default configuration that can be used for any RADIUS server that you configure the switch to communicate with. The default configuration includes:

- Encryption type
- Preshared Key
- Timeout value
- Number of retransmission attempts
- Allowing the user to specify a RADIUS server at login

About the Default RADIUS Server Encryption Type and Preshared Key

You need to configure the RADIUS preshared key to authenticate the switch to the RADIUS server. The length of the key is restricted to 64 characters and can include any printable ASCII characters (white spaces are not allowed). You can configure a global key to be used for all RADIUS server configurations on the switch.

You can override this global key assignment by explicitly using the **key** option when configuring and individual RADIUS server.

Configuring the Default RADIUS Server Encryption Type and Preshared Key

To configure the default RADIUS server encryption type and preshared key using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select **RADIUS.**

You see the RADIUS configuration in the Information pane.

Step 2 Click the **Defaults** tab.

You see the RADIUS default settings shown in Figure 35-3.

Figure 35-3 RADIUS Default Settings

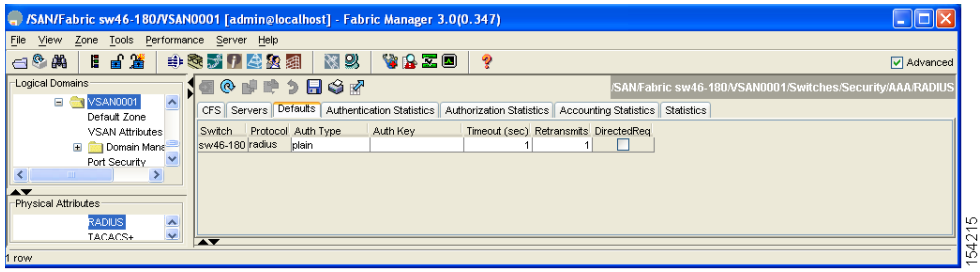

Step 3 Select **plain** or **encrypted** from the AuthType drop-down menu.

Step 4 Set the key in the Auth Key field.

Step 5 Click **Apply Changes** to save the changes.

Setting the Default RADIUS Server Timeout Interval and Retransmits

By default, a switch retries a RADIUS server only once. This number can be configured. The maximum is five retries per server. You can also configure the timeout value for the RADIUS server.

To configure the number of retransmissions and the time between retransmissions to the RADIUS servers using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA a**nd then select **RADIUS.**

You see the RADIUS configuration in the Information pane.

Step 2 Choose the **Defaults** tab.

You see the RADIUS default settings.

Step 3 Fill in the Timeout and Retransmits fields for authentication attempts.

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

About RADIUS Servers

You can add up to 64 RADIUS servers. RADIUS keys are always stored in encrypted form in persistent storage. The running configuration also displays encrypted keys. When you configure a new RADIUS server, you can use the default configuration or modify any of the parameters to override the default RADIUS configuration.

Configuring a RADIUS Server

To configure a RADIUS server and all its options using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select **RADIUS**.

You see the RADIUS configuration in the Information pane.

Step 2 Click the **Servers** tab.

You see any existing RADIUS servers.

Step 3 Click **Create Row** to add a new RADIUS server.

You see the Create RADIUS Server dialog box shown in Figure 35-4.

Figure 35-4 Create RADIUS Server

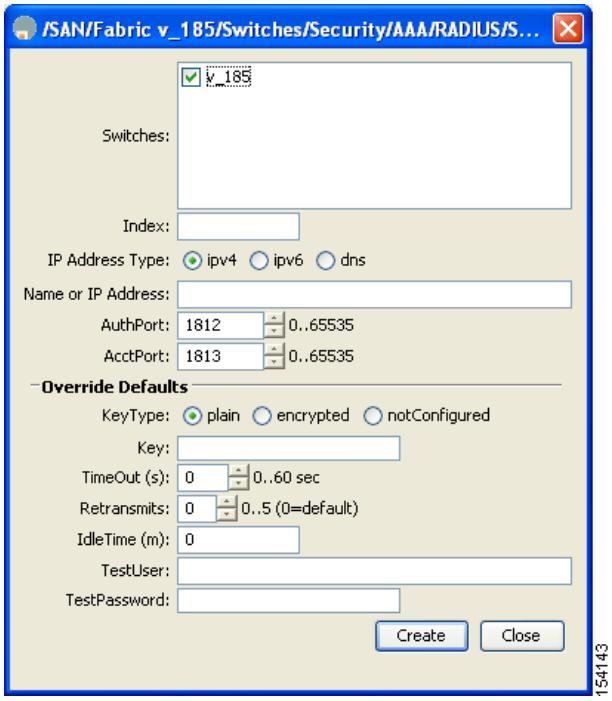

Step 4 Select the switches that you want to assign as RADIUS servers.

Step 5 Assign an index number to identify the RADIUS Server.

Step 6 Select the IP address type for the RADIUS server.

Step 7 Fill in the IP address or name for the RADIUS server.

Step 8 Optionally, modify the authentication and accounting ports used by this RADIUS server.

Step 9 Select the appropriate key type for the RADIUS server.

Step 10 Select the TimeOut value in seconds. The valid range is 0 to 60 seconds.

Step 11 Select the number of times the switch tries to connect to a RADIUS server(s) before reverting to local authentication.

Step 12 Enter the test idle time interval value in minutes. The valid range is 1 to 1440 minutes.

Step 13 Enter the test user with the default password. The default user name is test.

Step 14 Click **Create** to save these changes or click **Close** to discard any unsaved changes (see Figure 35-4).

Configuring the Test Idle Timer

The test idle timer specifies the interval during which a RADIUS server receives no requests before the MDS switch sends out a test packet.

Note The default idle timer value is 0 minutes. When the idle time interval is 0 minutes, periodic RADIUSserver monitoring is not performed.

To configure the test idle timer, see "Configuring a RADIUS Server" section on page 35-10.

Configuring Test User Name

You can configure a user name and password for periodic RADIUS server status testing. You do not need to configure the test user name and password to issue test messages to monitor RADIUS servers. You can use the default test user name (test) and default password (test).

Note We recommend that the test user name is not the same as an existing user in the RADIUS database for security reasons.

To configure the optional username and password for periodic RADIUS server status testing, see "Configuring a RADIUS Server" section on page 35-10.

About Validating a RADIUS Server

As of Cisco SAN-OS Release 3.0(1), you can periodically validate a RADIUS server. The switch sends a test authentication to the server using the user name and password that you configure. If the server does not respond to the test authentication, then the server is considered non-responding.

Note We recommend that, for security reasons, you do not use a user name that is configured on your RADIUS server as a test user name.

You can configure this option to test the server periodically, or you can run a one-time only test.

Periodically Validating a RADIUS Server

To configure the switch to periodically test a RADIUS server using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select **RADIUS**. You see the RADIUS configuration in the Information pane.

Step 2 Click the **Servers** tab.

You see any existing RADIUS servers.

Step 3 Click **Create Row** to add a new RADIUS server.

You see the Create RADIUS Server dialog box.

Step 4 Fill in the IP address.

Step 5 Optionally, modify the authentication and accounting ports used by this RADIUS server.

Step 6 Fill in the TestUser field and, optionally, the TestPassword field. The default password for the test is **Cisco**.

Step 7 Set the IdleTime field for the time that the server is idle before you send a test authentication.

Step 8 Click **Create** to save these changes or click **Close** to discard any unsaved changes.

Displaying RADIUS Server Statistics

To display RADIUS server statistics using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select **RADIUS**.

You see the RADIUS configuration in the Information pane.

Step 2 Click the **Statistics** tab.

You see the RADIUS server statistics.

About Users Specifying a RADIUS Server at Login

By default, a switch forwards an authentication request to the first server in the server group. As of Cisco SAN-OS Release 3.0(1), you can configure the switch to allow the user to specify which RADIUS server to authenticate with by enabling the directed request option. If you enable this option, the user can log in as *username@hostname*, where the *hostname* is the address of the RADIUS server.

Allowing Users to Specify a RADIUS Server at Login

To configure the switch to allow users to specify a RADIUS server at login using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select **RADIUS.**

You see the RADIUS configuration in the Information pane.

Step 2 Click the **Defaults** tab.

You see the RADIUS default settings.

Step 3 Check the **DirectedReq** check box for the RADIUS server.

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

About Vendor-Specific Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific attributes (VSAs) between the network access server and the RADIUS server. The IETF uses attribute 26. VSAs allow vendors to support their own extended attributes that are not suitable for general use. The Cisco RADIUS implementation supports one vendor-specific option using the format recommended in the specification. The Cisco vendor ID is 9, and the supported option is vendor type 1, which is named `cisco-avpair`. The value is a string with the following format:

```
protocol : attribute sep value *
```

Where `protocol` is a Cisco attribute for a particular type of authorization, and `sep` is `=` for mandatory attributes, and `*` is for optional attributes.

When you use RADIUS servers to authenticate yourself to a Cisco MDS 9000 Family switch, the RADIUS protocol directs the RADIUS server to return user attributes, such as authorization information, along with authentication results. This authorization information is specified through VSAs.

VSA Format

The following VSA protocol options are supported by the Cisco SAN-OS software:

- `Shell` protocol—used in access-accept packets to provide user profile information.
- `Accounting` protocol—used in accounting-request packets. If a value contains any white spaces, it should be put within double quotation marks.

The following attributes are supported by the Cisco SAN-OS software:

- `roles`–This attribute lists all the roles to which the user belongs. The value field is a string storing the list of group names delimited by white space. For example, if you belong to roles `vsan-admin` and `storage-admin`, the value field would be "`vsan-admin storage-admin`." This subattribute is sent in the VSA portion of the Access-Accept frames from the RADIUS server, and it can only be used with the shell protocol value. These are two examples using the roles attribute:

  ```
  shell:roles="network-admin vsan-admin"
  ```

  ```
  shell:roles*"network-admin vsan-admin"
  ```

  When an VSA is specified as `shell:roles*"network-admin vsan-admin"`, this VSA is flagged as an optional attribute, and other Cisco devices ignore this attribute.

- `accountinginfo`—This attribute stores additional accounting information besides the attributes covered by a standard RADIUS accounting protocol. This attribute is only sent in the VSA portion of the Account-Request frames from the RADIUS client on the switch, and it can only be used with the accounting protocol-related PDUs.

Specifying SNMPv3 on AAA Servers

The vendor/custom attribute `cisco-av-pair` can be used to specify user's role mapping using the format:

```
shell:roles="roleA roleB …"
```

The VSA format optionally specifies your SNMPv3 authentication and privacy protocol attributes also as follows:

```
shell:roles="roleA roleB..." snmpv3:auth=SHA priv=AES-128
```

The SNMPv3 authentication protocol options are SHA and MD5. The privacy protocol options are AES-128 and DES. If these options are not specified in the `cisco-av-pair` attribute on the ACS server, MD5 and DES are used by default.

Configuring TACACS+ Server Monitoring Parameters

A Cisco MDS switch uses the Terminal Access Controller Access Control System Plus (TACACS+) protocol to communicate with remote AAA servers. You can configure multiple TACACS+ servers and set timeout values.

TACACS+ is a client/server protocol that uses TCP (TCP port 49) for transport requirements. All switches in the Cisco MDS 9000 Family provide centralized authentication using the TACACS+ protocol. The addition of TACACS+ support in enables the following advantages over RADIUS authentication:

- Provides independent, modular AAA facilities. Authorization can be done without authentication.
- Uses the TCP transport protocol to send data between the AAA client and server, making reliable transfers with a connection-oriented protocol.
- Encrypts the entire protocol payload between the switch and the AAA server to ensure higher data confidentiality. The RADIUS protocol only encrypts passwords.

This section includes the following topics:

About TACACS+ Server Default Configuration

Fabric Manager allows you to set up a default configuration that can be used for any TACACS+ server that you configure the switch to communicate with. The default configuration includes:

- Encryption type
- Preshared Key
- Timeout value
- Number of retransmission attempts
- Allowing the user to specify a TACACS+ server at login

About the Default TACACS+ Server Encryption Type and Preshared Key

You need to configure the TACACS+ preshared key to authenticate the switch to the TACACS+ server. The length of the key is restricted to 64 characters and can include any printable ASCII characters (white spaces are not allowed). You can configure a global key to be used for all TACACS+ server configurations on the switch.

You can override this global key assignment by explicitly using the **key** option when configuring and individual TACACS+ server.

Setting the Default TACACS+ Server Encryption Type and Preshared Key

To configure the default TACACS+ server encryption type and preshared key using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select **TACACS+.**

You see the TACACS+ configuration in the Information pane.

Step 2 If the Defaults tab is greyed out, click the **CFS** tab.

Step 3 Click the **Defaults** tab.

You see the TACACS+ default settings.

Step 4 Select **plain** or **encrypted** from the AuthType drop-down menu and set the key in the Auth Key field.

Step 5 Click **Apply Changes** to save the changes.

Setting the Default TACACS+ Server Timeout Interval and Retransmits

By default, a switch retries a TACACS+ server only once. This number can be configured. The maximum is five retries per server. You can also configure the timeout value for the TACACS+ server.

To configure the number of retransmissions and the time between retransmissions to the TACACS+ servers using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select **TACACS+.**

You see the TACACS+ configuration in the Information pane.

Step 2 Choose the **Defaults** tab. (If the **Defaults** tab is disabled, click the **CFS** tab first.)

You see the TACACS+ default settings.

Step 3 Supply values for the Timeout and Retransmits fields for authentication attempts.

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

About TACACS+ Servers

By default, the TACACS+ feature is disabled in all switches in the Cisco MDS 9000 Family. Fabric Manager or Device Manager enables the TACACS+ feature automatically when you configure a TACACS+ server

If a secret key is not configured for a configured server, a warning message is issued if a global key is not configured. If a server key is not configured, the global key (if configured) is used for that server (see the "Setting the Default RADIUS Server Timeout Interval and Retransmits" section on page 35-9).

Note Prior to Cisco MDS SAN-OS Release 2.1(2), you can use the dollar sign ($) in the key but the key must be enclosed in double quotes, for example "k$". The percent sign (%) is not allowed. In Cisco MDS SAN-OS Release 2.1(2) and later, you can use the dollar sign ($) without double quotes and the percent sign (%) in global secret keys.

You can configure global values for the secret key for all TACACS+ servers.

Note If secret keys are configured for individual servers, those keys override the globally configured key.

Note Prior to Cisco MDS SAN-OS Release 2.1(2), you can use the dollar sign ($) in the key but the key must be enclosed in double quotes, for example "k$". The percent sign (%) is not allowed. In Cisco MDS SAN-OS Release 2.1(2) and later, you can use the dollar sign ($) without double quotes and the percent sign (%) in global secret keys.

Configuring a TACACS+ Server

To configure a TACACS+ server and all its options using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select **TACACS+**.

You see the TACACS+ configuration in the Information pane.

Step 2 Choose the **Servers** tab.

You see any existing TACACS+ servers.

Step 3 Click **Create Row** to add a new TACACS+ server.

You see the Create TACACS+ Server dialog box (see Figure 35-5).

Figure 35-5 Create TACACS+ Server Dialog Box

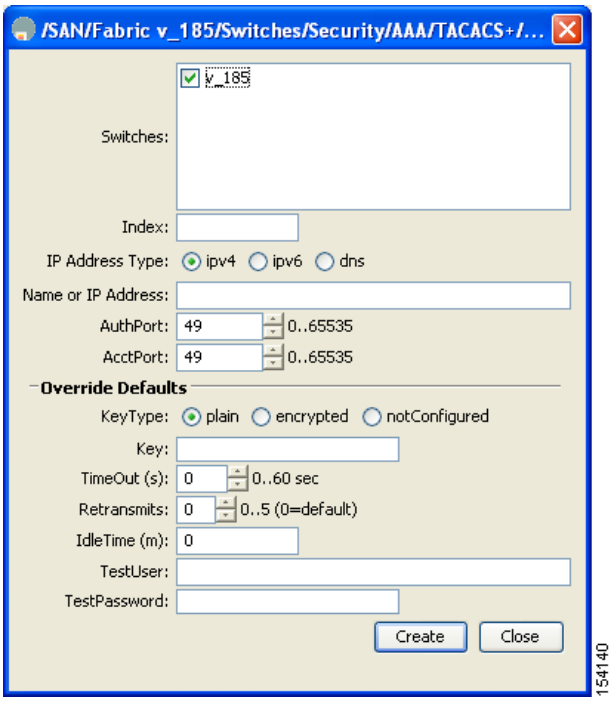

Step 4 Select the switches that you want to assign as TACACS servers.

Step 5 Assign an index number to identify the TACACS server.

Step 6 Select the IP address type for the TACACS server.

Step 7 Fill in the IP address or name for the TACACS server.

Step 8 Optionally, modify the authentication and accounting ports used by this TACACS server.

Step 9 Select the appropriate key type for the TACACS server.

Step 10 Select the TimeOut value in seconds. The valid range is 0 to 60 seconds.

Step 11 Select the number of times the switch tries to connect to a TACACS server(s) before reverting to local authentication.

Step 12 Enter the test idle time interval value in minutes. The valid range is 1 to 1440 minutes.

Step 13 Enter the test user with the default password. The default user name is test.

Step 14 Click **Create** to save these changes or click **Close** to discard any unsaved changes.

About Validating a TACACS+ Server

As of Cisco SAN-OS Release 3.0(1), you can periodically validate a TACACS+ server. The switch sends a test authentication to the server using the test user name and test password that you configure. If the server does not respond to the test authentication, then the server is considered non-responding.

Note We recommend that you do not configure the test user on your TACACS+ server for security reasons.

You can configure this option to test the server periodically, or you can run a one-time only test.

Periodically Validating a TACACS+ Server

To configure the switch to periodically test a TACACS+ server using Fabric Manager, see "Configuring a TACACS+ Server" section on page 35-16.

Displaying TACACS+ Server Statistics

To display TACACS+ server statistics using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select **TACACS+**.

You see the TACACS+ configuration in the Information pane.

Step 2 Choose the **Statistics** tab.

You see the TACACS+ server statistics.

About Users Specifying a TACACS+ Server at Login

By default, a switch forwards an authentication request to the first server in the server group. As of Cisco SAN-OS Release 3.0(1), you can configure the switch to allow the user to specify which TACACS+ server to authenticate with by enabling the directed request option. If you enable this option, the user can log in as username@hostname, where the hostname is the address of the TACACS+ server.

Allowing Users to Specify a TACACS+ Server at Login

To configure the switch to allow users to specify a TACACS+ server at login using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select **TACACS+.**

You see the TACACS+ configuration in the Information pane.

Step 2 Click the **Defaults** tab.

You see the TACACS+ default settings.

Step 3 Check the **DirectedReq** check box.

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

About Custom Attributes for Roles

Cisco MDS 9000 Family switches use the TACACS+ custom attribute for service shells to configure roles to which a user belongs. TACACS+ attributes are specified in `name=value` format. The attribute name for this custom attribute is `cisco-av-pair`. The following example illustrates how to specify roles using this attribute:

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

You can also configure optional custom attributes to avoid conflicts with non-MDS Cisco switches using the same AAA servers.

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

Additional custom attribute shell:roles are also supported:

```
shell:roles="network-admin vsan-admin"
```

or

```
shell:roles*"network-admin vsan-admin"
```

Note TACACS+ custom attributes can be defined on an Access Control Server (ACS) for various services (for example, shell). Cisco MDS 9000 Family switches require the TACACS+ custom attribute for the service shell to be used for defining roles.

Supported TACACS+ Servers

The Cisco SAN-OS software currently supports the following parameters for the listed TACACS+ servers:

- TACACS+

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

- Cisco ACS TACACS+

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

- Open TACACS+

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

Server Groups

You can specify one or more remote AAA servers to authenticate users using server groups. All members of a group must belong to the same protocol, either RADIUS or TACACS+. The servers are tried in the same order in which you configure them.

This section includes the following topics:

About Configuring Server Groups

You can configure these server groups at any time but they only take effect when you apply them to an AAA service. You configure AAA policies for CLI users or Fabric Manager or Device Manager users.

Configuring Server Groups

To configure a RADIUS or TACACS+ server group using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **AAA**.

You see the AAA configuration in the Information pane shown in Figure 35-6.

Figure 35-6 AAA Server Groups

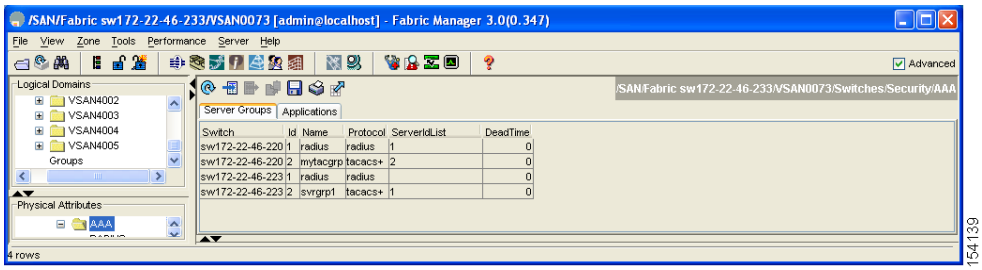

Step 2 If you don't see the screen in Figure 35-6, click the **Server Group** tab.

You see the RADIUS or TACACS+ server groups configured.

Step 3 Click **Create Row** to create a server group.

You see the Create Server dialog box.

Step 4 Select the **radius** radio button to add a RADIUS server group or select **tacacs+** to add a TACACS+ server group.

Step 5 Supply server names for the ServerIdList field.

Step 6 Set the DeadTime field for the number of minutes that a server can be nonresponsive before it is marked as bypassed. See the "About Bypassing a Nonresponsive Server" section on page 35-20.

Step 7 Click **Create** to create this server group or click **Close** to exit the dialog box without creating the new server group.

About Bypassing a Nonresponsive Server

As of Cisco SAN-OS Release 3.0(1), you can bypass a nonresponsive AAA server within a server group. If the switch detects a non-responsive server, it will bypass that server when authenticating users. Use this feature to minimize login delays caused by a faulty server. Instead of sending a request to a

non-responsive server and waiting for the authentication request to timeout, the switch sends the authentication request to the next server in the server group. If there are no other responding servers in the server group, the switch continues to attempt authentications against the non-responsive server.

AAA Server Distribution

Configuration for RADIUS and TACACS+ AAA on an MDS switch can be distributed using the Cisco Fabric Services (CFS). The distribution is disabled by default (see Chapter 12, "Using the CFS Infrastructure").

After enabling the distribution, the first server or global configuration starts an implicit session. All server configuration commands entered thereafter are stored in a temporary database and applied to all switches in the fabric (including the originating one) when you explicitly commit the database. The various server and global parameters are distributed, except the server and global keys. These keys are unique secrets to a switch and should not be shared with other switches.

Note Server group configurations are not distributed.

This section includes the following topics:

Enabling AAA Server Distribution

Only switches with distribution enabled can participate in the distribution activity.

To enable RADIUS server distribution using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select **RADIUS**.

You see the RADIUS configuration in the Information pane.

Step 2 Click the **CFS** tab. You see the RADIUS CFS configuration.

Step 3 Choose **enable** from the Admin drop-down list for all switches that you want to enable CFS for RADIUS (see Figure 35-7).

Figure 35-7 Enable a Radius Server on a Switch

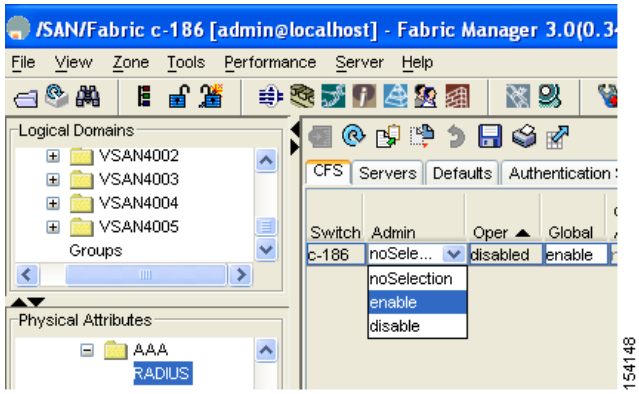

Step 4 Click **Apply Changes** to distribute these changes through the fabric.

To enable TACACS+ server distribution using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select **TACACS+**.

You see the TACACS+ configuration in the Information pane.

Step 2 Choose the **CFS** tab.

You see the TACACS+ CFS configuration.

Step 3 Choose **enable** from the Admin drop-down list for all switches that you want to enable CFS on for TACACS+.

Figure 35-8 Enable CFS TACACS+

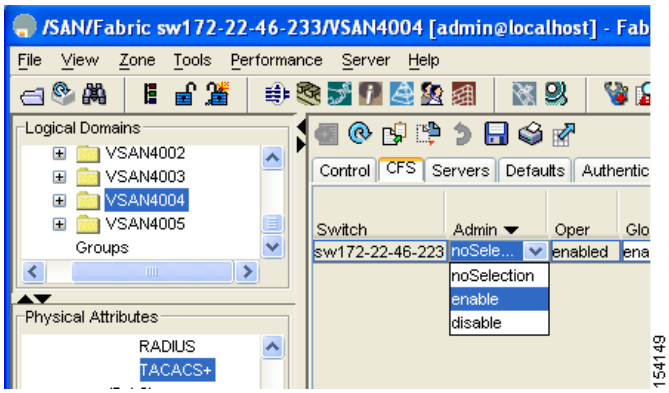

Step 4 Click **Apply Changes** to distribute these changes through the fabric.

Starting a Distribution Session on a Switch

A distribution session starts the moment you begin a RADIUS/TACACS+ server or global configuration. For example, the following tasks start an implicit session:

- Specifying the global timeout for RADIUS servers.
- Specifying the global timeout for TACACS+ servers.

Note After you issue the first configuration command related to AAA servers, all server and global configurations made (including the configuration that caused the distribution session start) are stored in a temporary buffer—not in the running configuration.

Displaying the Session Status

Once the implicit distribution session has started, you can check session status from Fabric Manager by expanding **Switches > Security > AAA** and selecting **RADIUS** or **TACACS+.** You see the distribution status on the CFS tab.

Displaying the Configuration to be Distributed

To display the RADIUS or TACACS+ global and/or server configuration stored in the temporary buffer using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select **RADIUS** or select **TACACS+.**

Step 2 Click the CFS tab.

You see the distribution status on the CFS tab.

Step 3 Select **pending** or **running** from the View Config As drop-down menu.

Figure 35-9 Config View As

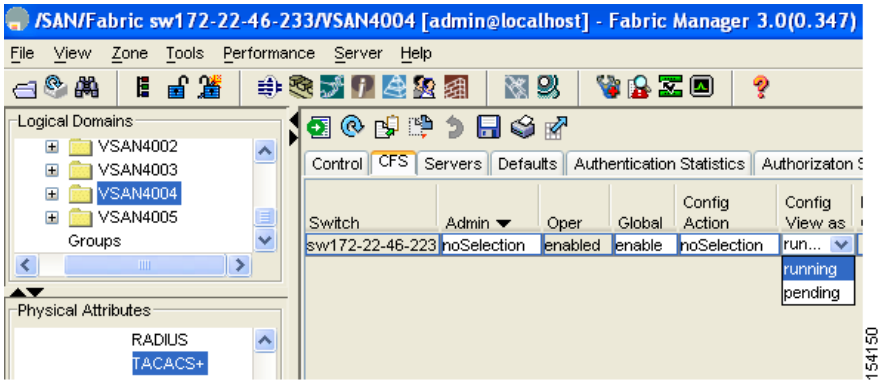

Step 4 Click **Apply Changes** to save the changes.

Step 5 Click the **Servers** tab to view the pending or running configuration.

Committing the Distribution

The RADIUS or TACACS+ global and/or server configuration stored in the temporary buffer can be applied to the running configuration across all switches in the fabric (including the originating switch).

To distribute a RADIUS or TACACS+ configuration using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select either **RADIUS** or **TACACS+**. You see the RADIUS or TACACS+ configuration in the Information pane.

Step 2 Choose the **CFS** tab. You see the RADIUS or TACACS+ CFS configuration.

Step 3 Choose **commit** in the Config Action drop-down list shown in Figure 35-10 for all switches that you want to enable CFS for RADIUS or TACACS+.

Figure 35-10 Committing the Configured Action

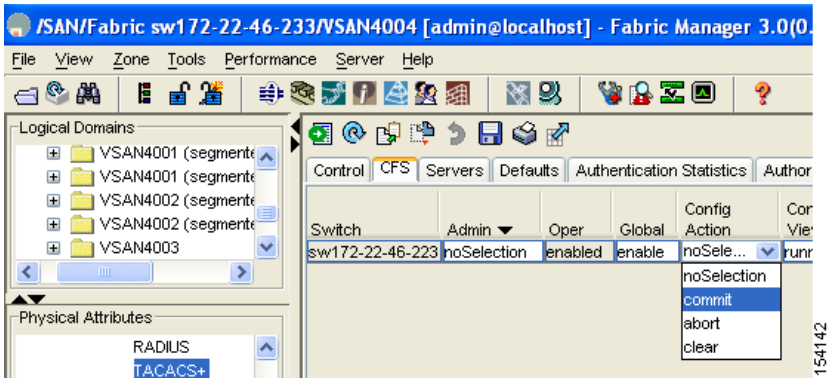

Step 4 Click **Apply Changes** to distribute the changes through the fabric.

Discarding the Distribution Session

Discarding the distribution of a session in progress causes the configuration in the temporary buffer to be dropped. The distribution is no applied.

To discard a RADIUS or TACACS+ distribution using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select either **RADIUS** or **TACACS+**. You see either the RADIUS or TACACS+ configuration in the Information pane.

Step 2 Click the **CFS** tab. You see either the RADIUS or TACACS+ CFS configuration.

Step 3 Choose **abort** from the Config Action drop-down list for each switch that should discard the pending RADIUS or TACACS+ distribution (see Figure 35-11).

Figure 35-11 Discarding a Pending TACACS+ Configuration

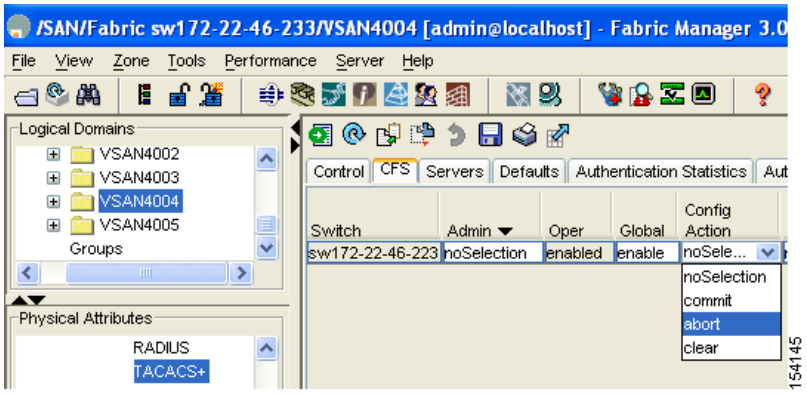

Step 4 Click **Apply Changes**.

Clearing Sessions

To clear a RADIUS or TACACS+ distribution using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select either **RADIUS** or **TACACS+**.

You see either the RADIUS or TACACS+ configuration in the Information pane.

Step 2 Choose the **CFS** tab. You see either the RADIUS or TACACS+ CFS configuration.

Step 3 Choose **clear** from the Config Action drop-down list for each switch that should clear the pending RADIUS or TACACS+ distribution (see Figure 35-12).

Figure 35-12 Clearing a TACACS+ Pending Distribution

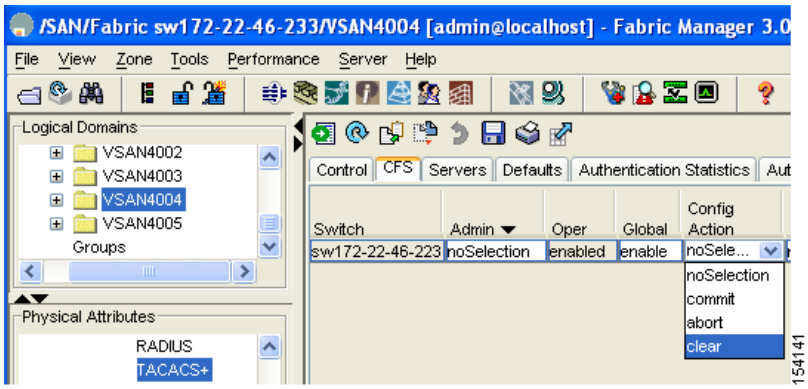

Step 4 Click **Apply Changes**.

Merge Guidelines for RADIUS and TACACS+ Configurations

The RADIUS and TACACS+ server and global configuration are merged when two fabrics merge. The merged configuration is applied to CFS distribution-enabled switches.

When merging the fabric, be aware of the following conditions:

- The server groups are not merged.
- The server and global keys are not changed during the merge.
- The merged configuration contains all servers found on all CFS enabled switches.
- The timeout and retransmit parameters of the merged configuration are the largest values found per server and global configuration.

Caution If there is a conflict between two switches in the server ports configured, the merge fails.

MSCHAP Authentication

Microsoft Challenge Handshake Authentication Protocol (MSCHAP) is the Microsoft version of CHAP. As of Cisco SAN-OS Release 3.0(1), you can use MSCHAP for user logins to a Cisco SAN-OS switch through a remote authentication server (RADIUS or TACACS+).

About Enabling MSCHAP

By default, the switch uses Password Authentication Protocol (PAP) authentication between the switch and the remote server. If you enable MSCHAP, you need to configure your RADIUS server to recognize the MSCHAP vendor-specific attributes. See the "About Vendor-Specific Attributes" section on page 35-13. Table 35-1 shows the RADIUS vendor-specific attributes required for MSCHAP.

Table 35-1 MSCHAP RADIUS Vendor-Specific Attributes

Vendor-ID Number	Vendor-Type Number	Vendor-Specific Attribute	Description
311	11	MSCHAP-Challenge	Contains the challenge sent by an AAA server to an MSCHAP user. It can be used in both Access-Request and Access-Challenge packets.
211	11	MSCHAP-Response	Contains the response value provided by an MS-CHAP user in response to the challenge. It is only used in Access-Request packets.

Enabling MSCHAP Authentication

To enable MSCHAP authentication using Device Manager, follow these steps:

Step 1 Click **Security > AAA.**

You see the AAA configuration in the Information pane.

Figure 35-13 AAA Configuration in Device Manager

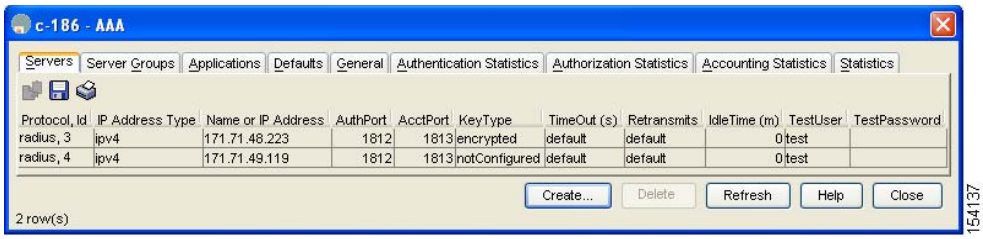

Step 2 Click the **General** tab.

You see the MSCHAP configuration shown in Figure 35-14.

Figure 35-14 MSCHAP Configuration

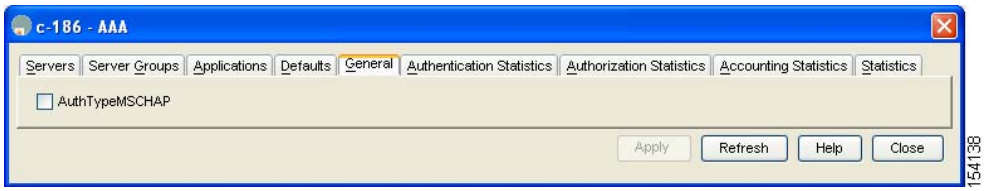

Step 3 Check the **AuthTypeMSCHAP** check box (see Figure 35-14) to use MSCHAP to authenticate users on the switch.

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

Local AAA Services

The system maintains the user name and password locally and stores the password information in encrypted form. You are authenticated based on the locally stored user information. See the "Configuring Users" section on page 33-12.

You can turn off password verification using the **none** option. If you configure this option, users can log in without giving a valid password. But the user should at least exist locally on the Cisco MDS 9000 Family switch.

Caution Use this option cautiously. If configured, any user can access the switch at any time.

Refer to the *Cisco MDS 9000 Family CLI Configuration Guide* to configure this option.

Default Settings

Table 35-2 lists the default settings for all switch security features in any switch.

Table 35-2 Default Switch Security Settings

Parameters	Default
AAA configuration services	Local.
Authentication port	1821.
Accounting port	1813.
Preshared key communication	Clear text.
RADIUS server time out	1 (one) second.
RADIUS server retries	Once.
TACACS+	Disabled.
TACACS+ servers	None configured.
TACACS+ server timeout	5 seconds.
AAA server distribution	Disabled.
VSAN policy for roles	Permit.
Accounting log size	250 KB.

CHAPTER 36

Configuring IPv4 Access Control Lists

Cisco MDS 9000 Family switches can route IP version 4 (IPv4) traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature routes traffic between VSANs. To do so, each VSAN must be in a different IP subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMS):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding on the in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

Note For information on IP version 6 (IPv6) support on the Cisco MDS 9000 Family switches, see

Switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. VRRP is a restartable application that provides a redundant, alternate path to the gateway switch.

IP Access Control Lists (IPv4-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IPv4-ACLs restrict IP-related traffic based on the configured IP filters. A filter contains the rules to match an IP packet, and if the packet matches, the rule also stipulates if the packet should be permitted or denied.

Each switch in the Cisco MDS 9000 Family can have a maximum of 64 IPv4-ACLs, and each IPv4-ACL can have a maximum of 256 filters.

This chapter includes the following sections:

IPv4-ACL Configuration Guidelines

Follow these guidelines when configuring IPv4-ACLs in any switch or director in the Cisco MDS 9000 Family:

- You could apply IPv4-ACLs to VSAN interfaces, the management interface, Gigabit Ethernet interfaces on IPS modules and MPS-14/2 modules, and Ethernet PortChannel interfaces.

Tip If IPv4-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group.

Caution Do not apply IPv4-ACLs to only one member of a PortChannel group. Apply IPv4-ACLs to the entire channel group.

- Configure the order of conditions accurately. As the IPv4-ACL filters are sequentially applied to the IP flows, only the first match determines the action taken. Subsequent matches are not considered. Be sure to configure the most important condition first. If no conditions match, the software drops the packet.

About Filter Contents

An IP filter contains rules for matching an IP packet based on the protocol, address, port, ICMP type, and type of service (TOS).

Protocol Information

The protocol information is required in each filter. It identifies the name or number of an IP protocol. You can specify the IP protocol in one of two ways:

- Specify an integer ranging from 0 to 255. This number represents the IP protocol.
- Specify the name of a protocol including, but not restricted to, Internet Protocol (IP), Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Control Message Protocol (ICMP).

Note When configuring IPv4-ACLs on Gigabit Ethernet interfaces, only use the TCP or ICMP options.

Address Information

The address information is required in each filter. It identifies the following details:

- Source: The address of the network or host from which the packet is being sent.
- Source-wildcard: The wildcard bits applied to the source.
- Destination: The number of the network or host to which the packet is being sent.
- Destination-wildcard: The wildcard bits applied to the destination.

Specify the source and source-wildcard or the destination and destination-wildcard in one of two ways:

- Using the 32-bit quantity in four-part, dotted decimal format (10.1.1.2/0.0.0.0 is the same as host 10.1.1.2).
  - Each wildcard bit set to zero indicates that the corresponding bit position in the packet's ip address must exactly match the bit value in the corresponding bit position in the source.
  - Each wildcard bit set to one indicates that both a zero bit and a one bit in the corresponding position of the packet's ip address will be considered a match to this access list entry. Place ones in the bit positions you want to ignore. For example, 0.0.255.255 requires an exact match of only the first 16 bits of the source. Wildcard bits set to one do not need to be contiguous in the source-wildcard. For example, a source-wildcard of 0.255.0.64 would be valid.
- Using the **any** option as an abbreviation for a source and source-wildcard or destination and destination-wildcard (0.0.0.0/255.255.255.255)

Port Information

The port information is optional. To compare the source and destination ports, use the **eq** (equal) option, the **gt** (greater than) option, the **lt** (less than) option, or the **range** (range of ports) option. You can specify the port information in one of two ways:

- Specify the number of the port. Port numbers range from 0 to 65535. Table 36-1 displays the port numbers recognized by the Cisco SAN-OS software for associated TCP and UDP ports.
- Specify the name of a TCP or UDP port as follows:
  - TCP port names can only be used when filtering TCP.
  - UDP port names can only be used when filtering UDP.

Table 36-1 TCP and UDP Port Numbers

Protocol	Port	Number
UDP	dns	53
	tftp	69
	ntp	123
	radius accounting	1646 or 1813
	radius authentication	1645 or 1812
	snmp	161
	snmp-trap	162
	syslog	514

Table 36-1 TCP and UDP Port Numbers (continued)

Protocol	Port	Number
TCP **Note** If the TCP connection is already established, use the **established** option to find matches. A match occurs if the TCP datagram has the ACK, FIN, PSH, RST, or URG control bit set.	ftp	20
	ftp-data	21
	ssh	22
	telnet	23
	smtp	25
	tasacs-ds	65
	www	80
	sftp	115
	http	143
	wbem-http	5988
	wbem-https	5989

ICMP Information

IP packets can be filtered based on the following optional ICMP conditions:

- The icmp-type: The ICMP message type is a number from 0 to 255.
- The icmp-code: The ICMP message code is a number from 0 to 255.

Table 36-2 displays the value for each ICMP type.

Table 36-2 ICMP Type Value

ICMP Type[1]	Code
echo	8
echo-reply	0
destination unreachable	3
traceroute	30
time exceeded	11

1. ICMP redirect packets are always rejected.

TOS Information

IP packets can be filtered based on the following optional TOS conditions:

- The TOS level: The level is a number from 0 to 15.
- The TOS name: The name can be max-reliability, max-throughput, min-delay, min-monetary-cost, and normal.

Creating IPv4-Access Lists with the IP-ACL Wizard

To create an ordered list of IP filters in a named IPv4-ACL profile using the IPv4-ACL Wizard in Fabric Manager, follow these steps:

Step 1 Choose the **IP-ACL Wizard** icon from the Fabric Manager toolbar. You see the IP-ACL Wizard.

Figure 36-1 IP-ACL Wizard

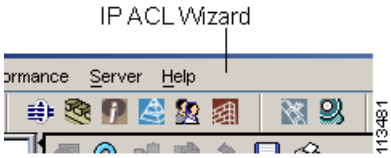

Step 2 Enter name for the IP-ACL.

Step 3 Click **Add** to add a new rule to this IP-ACL. You see a new rule in the table with default values.

Step 4 Modify the Source Ip and Source Mask as necessary for your filter.

Note The IP-ACL Wizard only creates inbound IP filters.

Step 5 Choose the appropriate filter type from the Application column.

Step 6 Choose **permit** or **deny** from the Action column.

Step 7 Repeat Step 3 through Step 6 for additional IP filters.

Step 8 Click **Up** or **Down** to order the filters in this IP-ACL.

Tip Order the IP filters carefully. Traffic is compared to the IP filters in order. The first match is applied and the rest are ignored.

Step 9 Click **Next**.

You see a list of switches that this IP-ACL can be applied to.

Step 10 Uncheck any switches that you do not want this IP-ACL applied to.

Step 11 Select the **Interface** you want this IP-ACL applied to.

Step 12 Click **Finish** to create this IP-ACL and apply it to the selected switches, or click **Cancel** to exit the IP-ACL Wizard without creating an IP-ACL.

IPv4-ACL Creation in Device Manager

Traffic coming into the switch is compared to IPv4-ACL filters based on the order that the filters occur in the switch. New filters are added to the end of the IPv4-ACL. The switch keeps looking until it has a match. If no matches are found when the switch reaches the end of the filter, the traffic is denied. For this reason, you should have the frequently hit filters at the top of the filter. There is an *implied deny* for traffic that is not permitted. A single-entry IPv4-ACL with only one deny entry has the effect of denying all traffic.

To configure an IPv4-ACL, you must complete the following tasks:

1. Create an IPv4-ACL by specifying a filter name and one or more access condition(s). Filters require the source and destination address to match a condition. Use optional keywords to configure finer granularity.
2. Apply the access filter to specified interfaces.

The IPv4-ACL Wizard in Fabric Manager provides tools to create an ordered list of simple IP filters and apply those filters to switches in the fabric.

This section includes the following topics:

Creating IPv4-ACLs in Device Manager

To create more complex IPv4-ACLs using Device Manager, follow these steps:

Step 1 Click **Security** and then select **IP ACLs**.

You see the IPv4-ACL dialog box Profiles tab shown in Figure 36-2.

Figure 36-2 IPv4-ACL Dialog Box Profiles

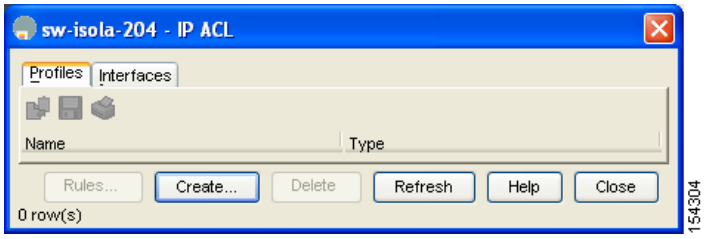

Step 2 Click **Create** to create an IPv4-ACL.

You see the Create IP ACL Profiles dialog box shown in Figure 36-3.

Figure 36-3 Create IP ACL Profiles

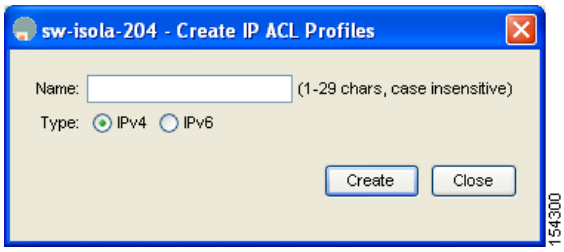

Step 3 Enter an IPv4-ACL profile name, click **Create,** and then click **Close**.

This creates a new empty IPv4-ACL profile.

Step 4 Click the IPv4-ACL you created and click **Rules**.

You see the list of IP filters associated with this IPv4-ACL (none for this new one) in Figure 36-4.

Figure 36-4 IP Filters Associated with the Current IPv4-ACL

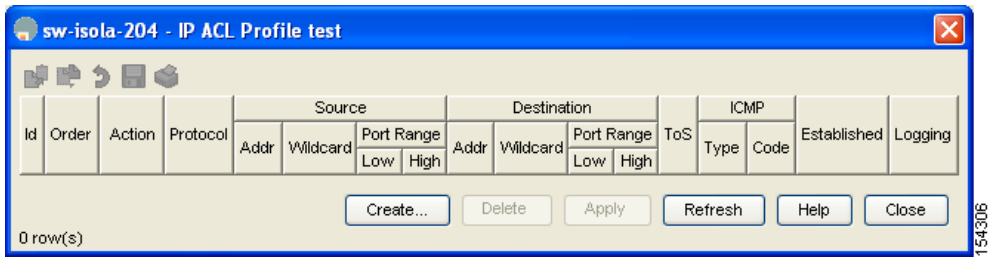

Step 5 Click **Create** to create an IP filter.

You see the Create IP Filter dialog box shown in Figure 36-5.

Figure 36-5 Create IP Filter Dialog Box

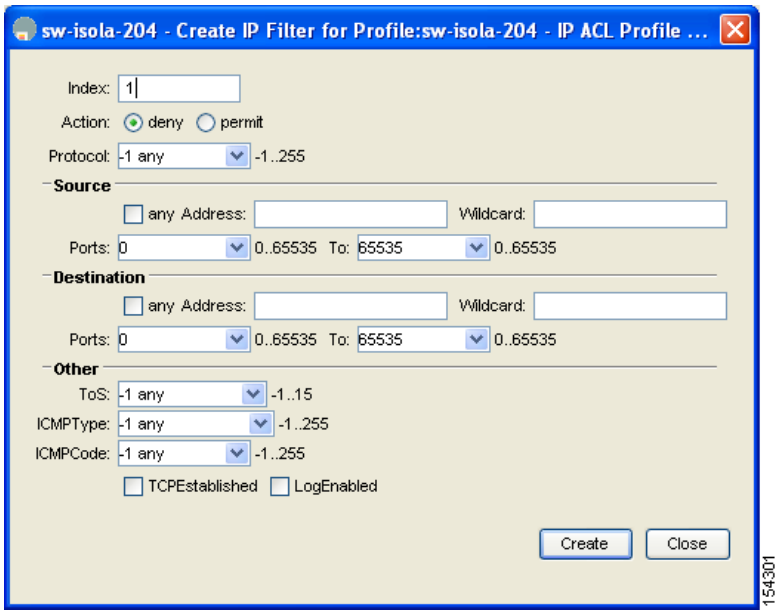

Step 6 Choose either **permit** or **deny** for the Action and set the IP Number in the Protocol field. The drop-down menu provides common filtered protocols.

Step 7 Set the source IP address you want this filter to match against and the wildcard mask, or check the **any** check box to match this filter against any IP address.

This creates an IP filter that will check the source IP address of frames.

Note The wildcard mask denotes a subset of the IP address you want to match against. This allows a range of addresses to match against this filter.

Step 8 Set the transport layer source port range if the protocol chosen is TCP or UDP.

Step 9 Repeat Step 7 and Step 8 for the destination IP address and port range.

This creates an IP filter that will check the destination IP address of frames.

Step 10 Set the ToS, ICMPType, and ICMPCode fields as appropriate.

Step 11 Check the **TCPEstablished** check box if you want to match TCP connections with ACK,FIN,PSH,RST,SYN or URG control bits set.

Step 12 Check the **LogEnabled** check box if you want to log all frames that match this IP filter.

Step 13 Click **Create** to create this IP filter and add it to your IPv4-ACL or click **Close** to close the IP Filter dialog box without creating an IP filter.

Any existing IP filters for this IPv4-ACL can be modified from the IPv4-ACL dialog box but the filters cannot be reordered.

Adding IP Filters to an Existing IPv4-ACL

After you create an IPv4-ACL, place subsequent addition IP filters at the end of the IPv4-ACL if you are using Device Manager. Fabric Manager allows you to reorder existing rules for a profile. Each configured entry is automatically added to the end of a IPv4-ACL.

To add entries to an existing IPv4-ACL using Device Manager, follow these steps:

Step 1 Click **Security** and then select **IP ACLs**.

You see the IP-ACL dialog box.

Step 2 Click the IP-ACL you want to modify and click **Rules**.

You see the list of IP filters associated with this IP-ACL in Figure 36-6.

Figure 36-6 IP Filters Applied to this IP-ACL

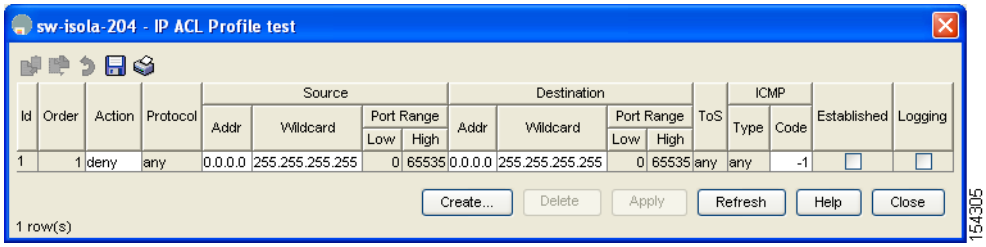

Step 3 Click **Create** to create an IP filter.

You see the Create IP Filter dialog box.

Step 4 Choose the **permit** or **deny** Action radio button and set the IP number in the Protocol field. The drop-down menu provides common filtered protocols.

Step 5 Set the source IP address you want this filter to match against and the wildcard mask, or check the **Any** check box to match this filter against any IP address. This creates an IP filter that will check the source IP address of frames.

Note The wildcard mask denotes a subset of the IP address you want to match against. This allows a range of addresses to match against this filter.

Step 6 Set the transport layer source port range if the protocol chosen is TCP or UDP.

Step 7 Repeat Step 7 and Step 8 for the destination IP address and port range. This creates an IP filter that will check the destination IP address of frames.

Step 8 Set the ToS, ICMPType, and ICMPCode fields as appropriate.

Step 9 Check the **TCPEstablished** check box if you want to match TCP connections with ACK,FIN,PSH,RST,SYN or URG control bits set.

Step 10 Check the **LogEnabled** check box if you want to log all frames that match this IP filter.

Step 11 Click **Create** to create this IP filter and add it to your IP-ACL or click **Close** to close the IP Filter dialog box without creating an IP filter.

Removing IP Filters from an Existing IPv4-ACL

To remove configured filters from an IPv4-ACL using Device Manager, follow these steps:

Step 1 Choose **Security > IP ACLs**.

You see the IP-ACL dialog box.

Step 2 Click the IP-ACL you want to modify and click **Rules**.

You see the list of IP filters associated with this IP-ACL.

Step 3 Select the filter that you want to delete and click **Delete** to delete that IP filter.

About Applying an IPv4-ACL to an Interface

You can define IPv4-ACLs without applying them. However, the IPv4-ACLs have no effect until they are applied to an interface on the switch.

Tip Apply the IPv4-ACL on the interface closest to the source of the traffic.

When you are trying to block traffic from source to destination, you can apply an inbound IPv4-ACL to M0 on Switch 1 instead of an outbound filter to M1 on Switch 3 (see Figure 36-7).

Figure 36-7 Denying Traffic on the Inbound Interface

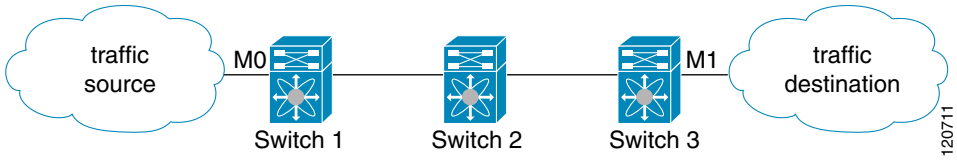

The **access-group** option controls access to an interface. Each interface can only be associated with one IPv4-ACL per direction. The ingress direction can have a different IPv4-ACL than the egress direction. The IPv4-ACL becomes active when applied to the interface.

Tip Create all conditions in an IPv4-ACL before applying it to the interface.

Caution If you apply an IPv4-ACL to an interface before creating it, all packets in that interface are dropped because the IPv4-ACL is empty.

The terms *in, out, source*, and *destination* are used as referenced by the switch.

- In—Traffic that arrives at the interface and goes through the switch; the source is where it transmitted from and the destination is where it is transmitted to (on the other side of the router).

Tip The IPv4-ACL applied to the interface for the ingress traffic affects both local and remote traffic.

- Out—Traffic that has already been through the switch and is leaving the interface; the source is where it transmitted from and the destination is where it is transmitted to.

Tip The IPv4-ACL applied to the interface for the egress traffic only affects local traffic.

Applying an IPv4-ACL to an Interface

To apply an IPv4-ACL to an interface using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **IP ACL** in the Physical Attributes pane.

You see the IP-ACL configuration in the Information pane.

Step 2 Click the **Interfaces** tab.

You see a list of interfaces and associated IP-ACLs.

Step 3 Click **Create Row**.

You see the Create Interface dialog box.

Step 4 Optionally, select the switches you want to include in the IP-ACL by checking the check boxes next to the switch addresses in Fabric Manager.

Step 5 Set the interface you want associated with an IP-ACL in the Interface field.

Step 6 Choose a ProfileDirection (either **inbound** or **outbound**).

Step 7 Enter the IP-ACL name in the Profile Name field.

Note This IP-ACL name must already have been created using the Create Profiles dialog box. If not, no filters will be enabled until you to go to the Create Profiles dialog box and create the profile.

Step 8 Click **Create** to associate the IP-ACL, or click **Close** to close the Create Interfaces dialog box without associating an access list.

You see the newly associated access list in the list of IP-ACLs.

Deleting IPv4-ACL

You must delete the association between the IPv4-ACL and interfaces before deleting the IPv4-ACL.

To delete an IPv4-ACL using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **IP ACL** from the Physical Attributes pane.

You see the IPv4-ACL configuration in the Information pane.

Step 2 Click the **Profiles** tab.

You see a list of switches, ACLs, and profile names.

Step 3 Select the row you want to delete. To delete multiple rows, hold down the Shift key while selecting rows.

Step 4 Click **Delete Row**. The IPv4-ACLs are deleted.

Reading the IPv4-ACL Log Dump

Use the LogEnabled check box option during IP filter creation to log information about packets that match this filter. The log output displays the ACL number, permit or deny status, and port information.

For the input ACL, the log displays the raw MAC information. The keyword “MAC=” does not refer to showing an Ethernet MAC frame with MAC address information. It refers to the Layer 2 MAC-layer information dumped to the log. For the output ACL, the raw Layer 2 information is not logged.

The following example is an input ACL log dump.

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN=vsan1 OUT=
MAC=10:00:00:05:30:00:47:df:10:00:00:05:30:00:8a:1f:aa:aa:03:00:00:00:08:00:45:00:00:54:00
:00:40:00:40:01:0e:86:0b:0b:0b:0c:0b:0b:0b:02:08:00:ff:9c:01:15:05:00:6f:09:17:3f:80:02:01
:00:08:09:0a:0b:0c:0d:0e:0f:10:11:12:13:14:15:16:17:18:19:1a:1b:1c:1d:1e:1f:20:21:22:23:24
:25:26:27:28:29:2a:2b SRC=11.11.11.12 DST=11.11.11.2 LEN=84 TOS=0x00 PREC=0x00 TTL=64 ID=0
DF PROTO=ICMP TYPE=8 CODE=0 ID=277 SEQ=1280
```

The following example is an output ACL log dump.

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN= OUT=vsan1 SRC=11.11.11.2 DST=11.11.11.12 LEN=84 TOS=0x00 PREC=0x00
TTL=255 ID=38095 PROTO=ICMP TYPE=0 CODE=0 ID=277 SEQ=1280
```

Example IP ACL Configuration

To define an IPv4-ACL that restricts management access using Device Manager, follow these steps:

Step 1 Expand **Security** and then select **IP ACLs**.

You see the IP-ACL dialog box in Figure 36-8.

Figure 36-8 IPv4-ACL Dialog Box Profiles Tab

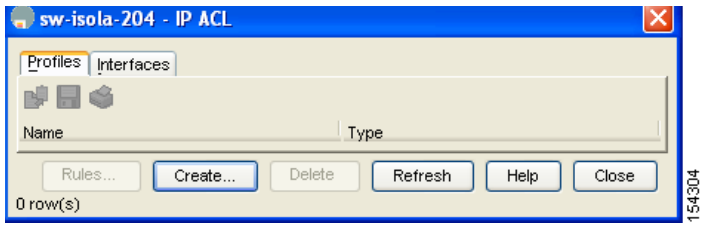

Step 2 Click **Create** to create an IP-ACL.

You see the Create IP ACL Profiles dialog box shown in Figure 36-9.

Figure 36-9 Create IP ACL Profiles Dialog Box

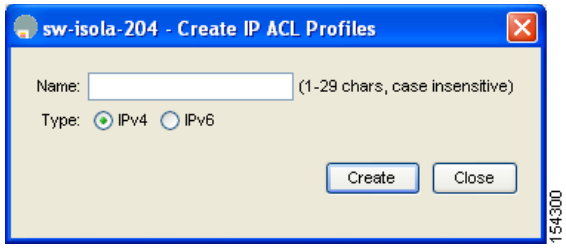

Step 3 Enter **RestrictMgmt** as the profile name and click **Create** (see Figure 36-9).

This creates an empty, IP-ACL named RestrictMgmt in Figure 36-10.

Figure 36-10 RestrictMgmt Profile Added to the List

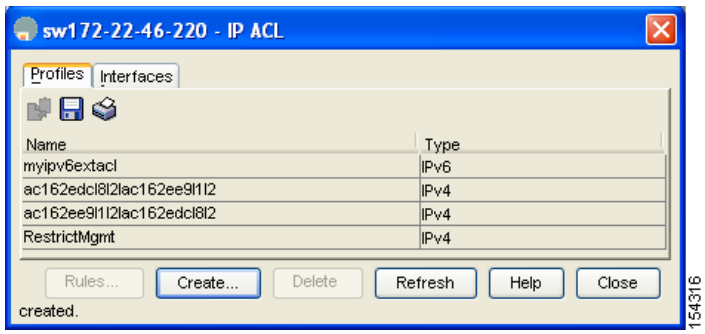

Step 4 Click **RestrictMgmt** and click **Rules** (see Figure 36-10).

You see an empty list of IP filters associated with this IP-ACL in Figure 36-11.

Figure 36-11 Empty list of IP Filters Associated with RestrictMgmt

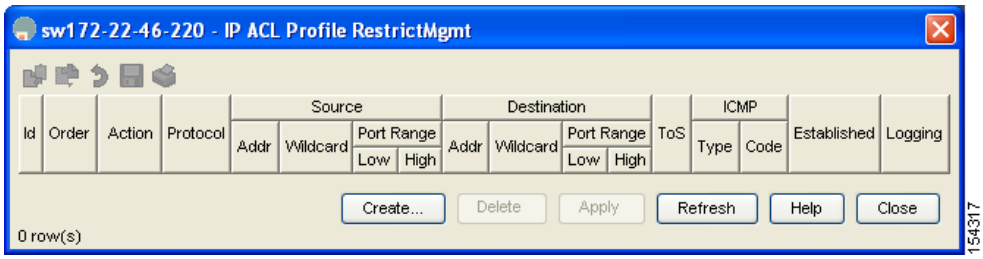

Step 5 Click **Create** to create the first IP filter (see Figure 36-11).

You see the Create IP Filter dialog box in Figure 36-12.

Figure 36-12 Create IP Filter Dialog Box

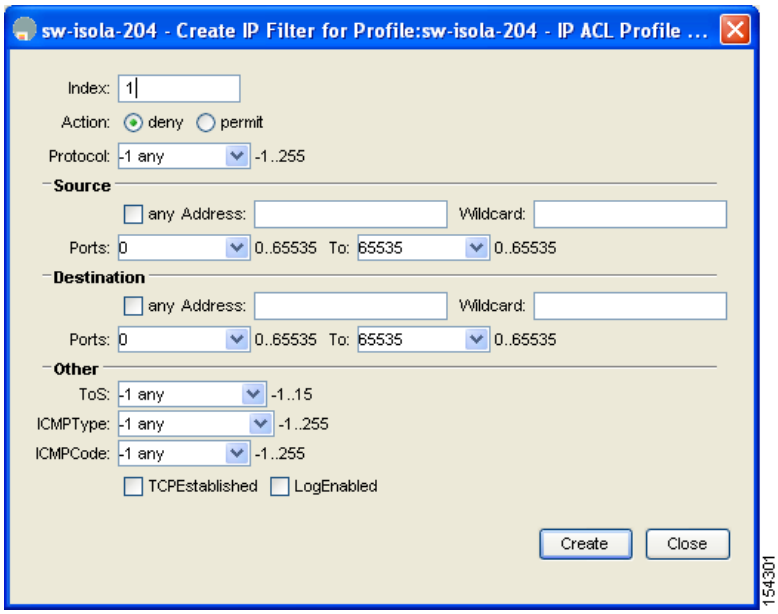

Step 6 Create an IP filter to allow management communications from a trusted subnet:

a. Choose the **permit** Action and select **0 IP** from the Protocol drop-down menu (see Figure 36-13).

b. Set the source IP address to **10.67.16.0** and the wildcard mask to **0.0.0.255** (see Figure 36-13).

Note The wildcard mask denotes a subset of the IP Address you want to match against. This allows a range of addresses to match against this filter.

c. Check the **any** check box for the destination address (see Figure 36-13).

Figure 36-13 IRestrict_Mgmt Filter Criteria

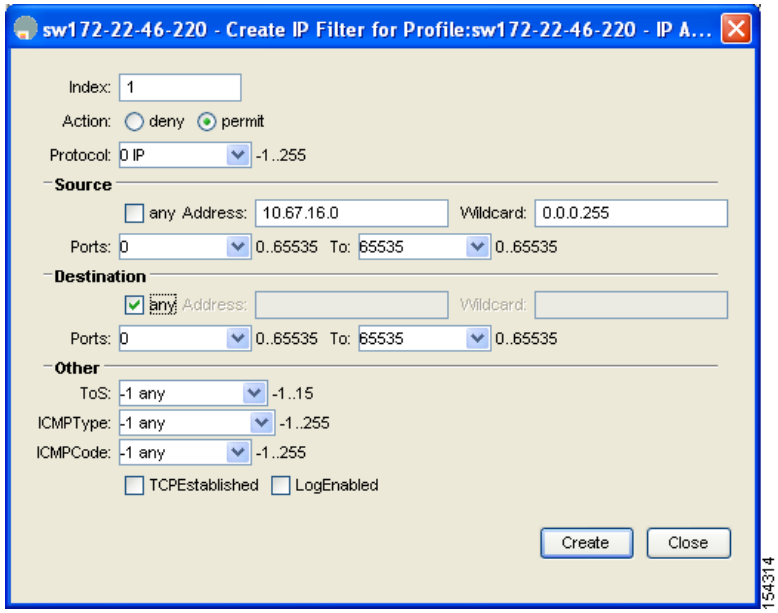

d. Click **Create** to create this IP filter and add it to the RestrictMgmt IP-ACL (see Figure 36-13).

Step a. through step d. create an IP filter that allows communications for all addresses in the 10.67.16.0/24 subnet.

Step 7 Create an IP filter to allow ICMP ping commands:

a. Choose the **permit** Action and select **1-ICMP** from the Protocol drop-down menu (see Figure 36-14).

b. Check the **any** check box for the source address.

c. Check the **any** check box for the destination address.

d. Select **8 echo** from the ICMPType drop-down menu.

Figure 36-14 Allow ICMP Ping Commands Filter Criteria

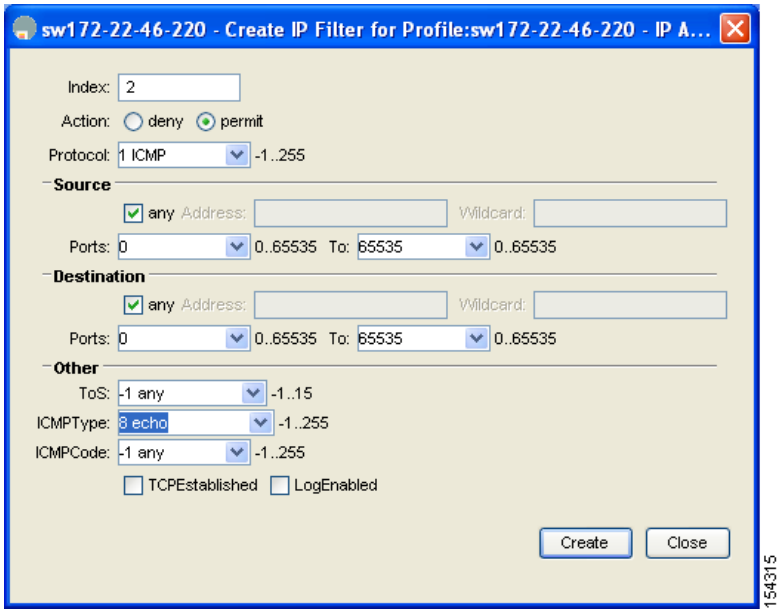

e. Click **Create** to create this IP filter and add it to the RestrictMgmt IP-ACL (see Figure 36-14).

Step a. through step e. create an IP filter that allows ICMP ping.

Step 8 Create a final IP Filter to block all other traffic:

a. Choose the **deny** Action and select **0 IP** from the Protocol drop-down menu (see Figure 36-15).

b. Check the **any** check box for the source address.

c. Check the **any** check box for the destination address.

Figure 36-15 Filter to Block All Other Traffic

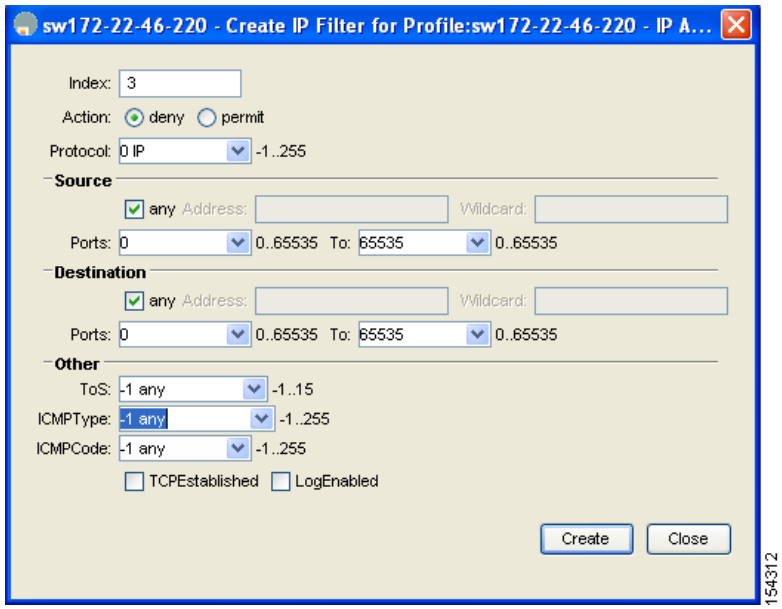

d. Click **Create** to create this IP filter and add it to the RestrictMgmt IP-ACL (see Figure 36-15).

e. Click **Close** to close the Create IP Filter dialog box.

Step a. through step d. create an IP filter that blocks all other traffic.

Step 9 Apply the RestrictMgmt IP ACL to the mgmt0 interface:

a. Click **Security,** select **IP ACL** and then click the **Interfaces** tab in the IP ACL dialog box.

b. Click **Create**.

You see the Create IP-ACL Interfaces dialog box shown in Figure 36-16.

Figure 36-16 Create IP-ACL Interfaces Dialog Box

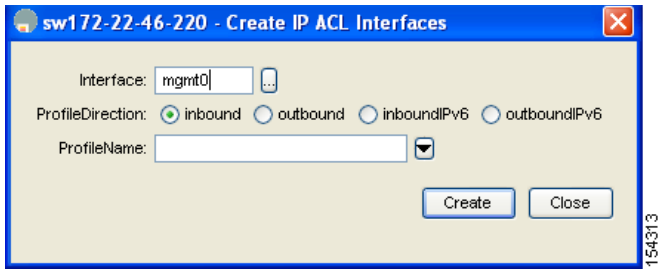

c. Select **mgmt0** from the Interfaces drop-down menu.

d. Select the **inbound** Profile Director.

e. Select **RestrictMgmt** from the ProfileName drop-down menu.

f. Click **Create** to apply the RestrictMgmt IP-ACL to the mgmt0 interface.

Step a. through step f. apply the new IP-ACL to the mgmt0 interface.

CHAPTER 37

Configuring IPv6 Access Control Lists

IP version 6 Access Control Lists (IPv6-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IPv6-ACLs restrict IP-related traffic based on the configured IP filters. A filter contains the rules to match an IP packet, and if the packet matches, the rule also stipulates if the packet should be permitted or denied.

Each switch in the Cisco MDS 9000 Family can have a maximum of 64 IPv6-ACLs and each IPv6-ACL can have a maximum of 256 filters.

This chapter includes the following sections:

IPv6-ACL Configuration Guidelines

Follow these guidelines when configuring IPv6-ACLs in any switch or director in the Cisco MDS 9000 Family:

- Apply IPv6-ACLs to VSAN interfaces, the management interface, Gigabit Ethernet interfaces on IPS modules and MPS-14/2 modules, and Ethernet PortChannel interfaces.

Tip If IPv6-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group. See the "Gigabit Ethernet IPv6-ACL Guidelines" section on page 48-23 for guidelines.

Caution Do not apply IPv6-ACLs to only one member of a PortChannel group. Apply IPv6-ACLs to the entire channel group.

- Configure the order of conditions accurately. As the IPv6-ACL filters are sequentially applied to the IP flows, only the first match determines the action taken. Subsequent matches are not considered. Be sure to configure the most important condition first. If no conditions match, the software drops the packet.

About Filter Contents

An IP filter contains rules for matching an IP packet based on the protocol, address, port, ICMP type, and type of service (TOS).

This section includes the following topics:

Protocol Information

Protocol information is required in each filter. It identifies the name or number of an IP protocol. You can specify the IP protocol in one of two ways:

- Specify an integer ranging from 0 to 255. This number represents the IP protocol.
- Specify the name of a protocol including, but not restricted to, Internet Protocol (IP), Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Control Message Protocol (ICMP).

Note When configuring IPv6-ACLs on Gigabit Ethernet interfaces, only use the TCP or ICMP options.

Address Information

Address information is required in each filter. It identifies the following details:

- Source: The address of the network or host from which the packet is being sent.
- Source-wildcard: The wildcard bits applied to the source.
- Destination: The number of the network or host to which the packet is being sent.
- Destination-wildcard: The wildcard bits applied to the destination.

Specify the source and source-wildcard or the destination and destination-wildcard in one of two ways:

- Use the 128-bit quantity in colon-separated hexadecimal format.
  - Each wildcard bit set to zero indicates that the corresponding bit position in the packet's IPv6 address must exactly match the bit value in the corresponding bit position in the source.
  - Each wildcard bit set to one indicates that both a zero bit and a one bit in the corresponding position of the packet's IPv6 address will be considered a match to this access list entry. Place ones in the bit positions you want to ignore. For example, 2001:0DB8:800:200C:/64 to require an exact match of only the first 16 bits of the source. Wildcard bits set to one do not need to be contiguous in the source-wildcard.
- Use the **any** option as an abbreviation for a source and source-wildcard or destination and destination-wildcard.

Port Information

Port information is optional. To compare the source and destination ports, use the **eq** (equal) option, the **gt** (greater than) option, the **lt** (less than) option, or the **range** (range of ports) option. You can specify the port information in one of two ways:

- Specify the number of the port. Port numbers range from 0 to 65535. Table 37-1 displays the port numbers recognized by the Cisco SAN-OS software for associated TCP and UDP ports.
- Specify the name of a TCP or UDP port as follows:
  - TCP port names can only be used when filtering TCP.
  - UDP port names can only be used when filtering UDP.

Table 37-1 TCP and UDP Port Numbers

Protocol	Port	Number
UDP	dns	53
	tftp	69
	ntp	123
	radius accounting	1646 or 1813
	radius authentication	1645 or 1812
	snmp	161
	snmp-trap	162
	syslog	514
TCP[1]	ftp	20
	ftp-data	21
	ssh	22
	telnet	23
	smtp	25
	tasacs-ds	65
	www	80
	sftp	115
	http	143
	wbem-http	5988
	wbem-https	5989

1. If the TCP connection is already **established**, use the **established** option to find matches. A match occurs if the TCP datagram has the ACK, FIN, PSH, RST, or URG control bit set.

ICMP Information

IP packets can be filtered based on the following optional ICMP conditions:

- The icmp-type: The ICMP message type is a number from 0 to 255.
- The icmp-code: The ICMP message code is a number from 0 to 255.

Table 37-2 displays the value for each ICMP type.

Table 37-2 ICMP Type Value

ICMP Type[1]	Code
echo	8
echo-reply	0
destination unreachable	3
traceroute	30
time exceeded	11

1. ICMP redirect packets are always rejected.

TOS Information

IP packets can be filtered based on the following optional TOS conditions:

- The TOS level: The level is a number from 0 to 15.
- The TOS name: The name can be max-reliability, max-throughput, min-delay, min-monetary-cost, and normal.

Creating IPv6-ACLs with the IP-ACL Wizard

Traffic coming into the switch is compared to IPv6-ACL filters based on the order that the filters occur in the switch. New filters are added to the end of the IPv6-ACL. The switch keeps looking until it has a match. If no matches are found when the switch reaches the end of the filter, the traffic is denied. For this reason, you should have the frequently hit filters at the top of the filter. There is an *implied deny* for traffic that is not permitted. A single-entry IPv6-ACL with only one deny entry has the effect of denying all traffic.

To configure an IPv6-ACL, you must complete the following tasks:

Step 1 Create an IPv6-ACL by specifying a filter name.

Step 2 Add entries that contain the required source and destination addresses to match a condition. Use optional keywords to configure finer granularity.

Note The filter entries are executed in sequential order. You can only add the entries to the end of the list. Take care to add the entries in the correct order.

Step 3 Apply the access filter to specified interfaces.

See the "Gigabit Ethernet IPv6-ACL Guidelines" section on page 48-23.

Step 1 In Fabric Manager, choose the **IP-ACL Wizard** icon from the Fabric Manager toolbar. You see the IP-ACL Wizard.

Figure 37-1 IP-ACL Wizard

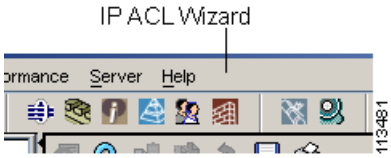

Step 2 Enter name for the IP-ACL.

Step 3 Click the extendedIPv6 option.

Step 4 Click the **Add** button to add a new rule to this IP-ACL. You see a new rule in the table with default values.

Step 5 Modify the Source IP and Source Mask as necessary for your filter.

Note The IP-ACL Wizard only creates inbound IP filters.

Step 6 Choose the appropriate filter type from the Application column.

Step 7 Choose **permit** or **deny** from the Action column.

Step 8 Repeat Step 4 through Step 7 for additional IP filters.

Step 9 Click **Up** or **Down** to order the filters in this IPv6-ACL.

Tip Order the IP filters carefully. Traffic is compared to the IP filters in order. The first match is applied and the rest are ignored.

Step 10 Click **Next**. You see a list of switches that this IPv6-ACL can be applied to.

Step 11 Uncheck any switches that you do not want this IPv6-ACL applied to.

Step 12 Select the **Interface** you want this IPv6-ACL applied to.

Step 13 Click **Finish** to create this IPv6-ACL and apply it to the selected switches, or click **Cancel** to exit the IPv6-ACL Wizard without creating an IP-ACL.

Adding Filters to an Existing IPv6-ACL

After you create an IPv6-ACL, you place subsequent additions at the end of the IPv6-ACL. You cannot insert filters in the middle of an IPv6-ACL. Each configured entry is automatically added to the end of an IPv6-ACL.

To add entries to an existing IPv6-ACL using Device Manager, follow these steps:

Step 1 Click **Security** and then select **IP ACLs**.

You see the IP ACL dialog box Profiles tab shown in Figure 37-2.

Figure 37-2 IPv6-ACL Dialog Box Profiles Tab

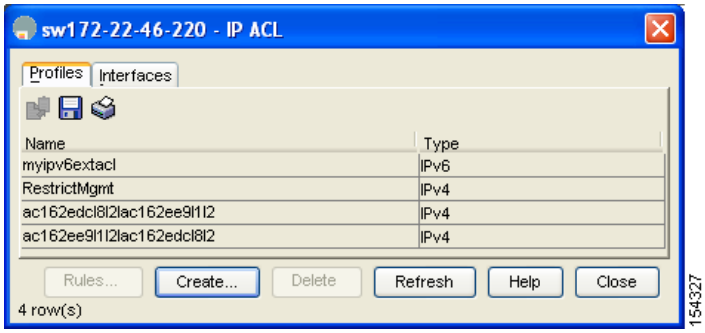

Step 2 Click the **IPv6** IP-ACL you want to modify (in Figure 37-2, the only IPv6 is **myipv6extacl**) and click **Rules**.

You see the list of IP filters associated with this IPv6 IP-ACL in Figure 37-3.

Figure 37-3 List of Filters for Myipv6extacl

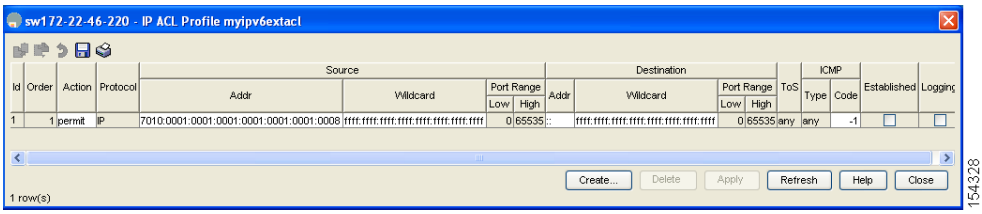

Step 3 Click **Create** to create another IP filter.

You see the Create IP Filter dialog box in Figure 37-4.

Figure 37-4 Create IP Filter Dialog Box

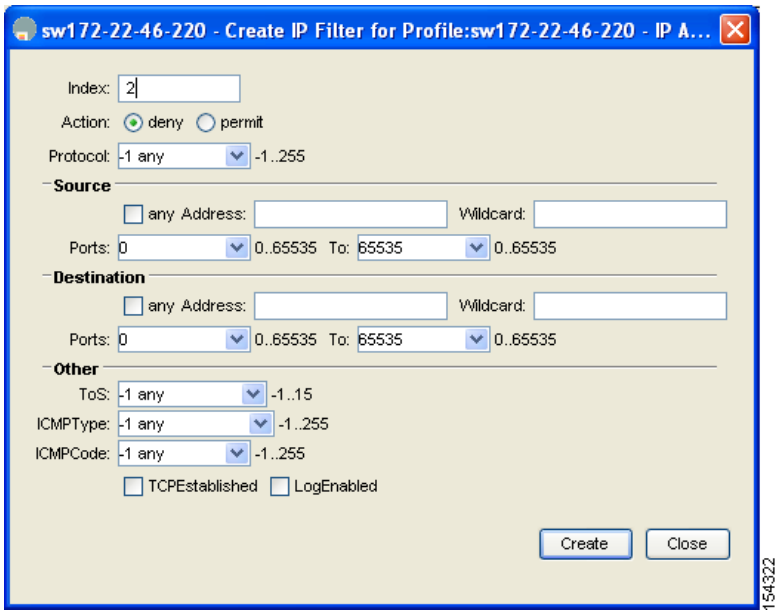

Step 4 Choose the **permit** or **deny** Action and set the Internet Protocol Number in the Protocol field. The drop-down menu provides common filtered protocols.

Step 5 Set the source IP address you want this filter to match against and the wildcard mask, or check the **any** check box to match this filter against any IP address.

This creates an IP filter that will check the source IP address of frames.

Note The wildcard mask denotes a subset of the IP Address you want to match against. This allows a range of addresses to match against this filter.

Step 6 Set the transport layer source port range if the protocol chosen is TCP or UDP.

Step 7 Repeat Step 6 and Step 7for the destination IP address and port range.

This creates an IP filter that will check the destination IP address of frames.

Step 8 Set ToS, ICMPType, and ICMPCode as appropriate.

Step 9 Check the **TCPEstablished** check box if you want to match TCP connections with ACK,FIN,PSH,RST,SYN or URG control bits set.

Step 10 Check the **LogEnabled** check box if you want to log all frames that match this IP filter.

Step 11 Click **Create** to create this IP Filter and add it to your IP-ACL or click **Close** to close the IP Filter dialog box without creating an IP filter.

Removing Entries from an Existing IPv6-ACL

To remove configured filters from an IPv6-ACL using Device Manager, follow these steps:

Step 1 Click **Security** and then select **IP ACLs**.

You see the IP ACL dialog box Profiles tab.

Step 2 Click the **IPv6** IP-ACL you want to modify (in Figure 37-5, the only IPv6 is **myipv6extacl**) and click **Rules**.

You see the list of IP filters associated with this IPv6 IP-ACL in Figure 37-5.

Figure 37-5 List of Filters for Myipv6extacl

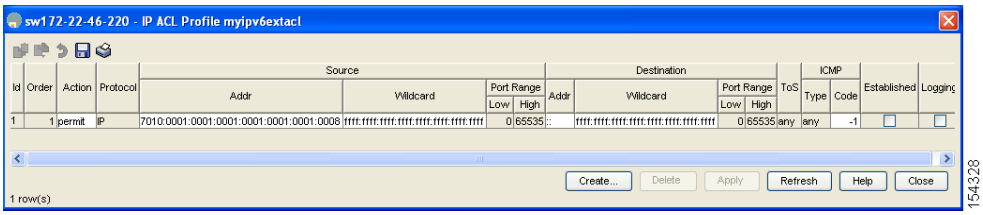

Step 3 Click the filter that you want to delete and click **Delete** to delete that IP filter.

Reading the IPv6-ACL Log Dump

Use the **log-deny** option at the end of a filter condition to log information about packets that match dropped entries. The log output displays the ACL number, permit or deny status, and port information.

For the input ACL, the log displays the raw MAC information. The keyword “MAC=” does not refer to showing an Ethernet MAC frame with MAC address information. It refers to the Layer 2 MAC-layer information dumped to the log. For the output ACL, the raw Layer 2 information is not logged.

The following example shows an input ACL log dump.

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN=vsan1 OUT=
MAC=10:00:00:05:30:00:47:df:10:00:00:05:30:00:8a:1f:aa:aa:03:00:00:00:08:00:45:00:00:54:00
:00:40:00:40:01:0e:86:0b:0b:0b:0c:0b:0b:0b:02:08:00:ff:9c:01:15:05:00:6f:09:17:3f:80:02:01
:00:08:09:0a:0b:0c:0d:0e:0f:10:11:12:13:14:15:16:17:18:19:1a:1b:1c:1d:1e:1f:20:21:22:23:24
:25:26:27:28:29:2a:2b SRC=11.11.11.12 DST=11.11.11.2 LEN=84 TOS=0x00 PREC=0x00 TTL=64 ID=0
DF PROTO=ICMP TYPE=8 CODE=0 ID=277 SEQ=1280
```

The following example is an output ACL log dump.

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN= OUT=vsan1 SRC=11.11.11.2 DST=11.11.11.12 LEN=84 TOS=0x00 PREC=0x00
TTL=255 ID=38095 PROTO=ICMP TYPE=0 CODE=0 ID=277 SEQ=1280
```

Applying an IPv6-ACL to an Interface

You can define IPv6-ACLs without applying them. However, the IPv6-ACLs will have no effect until they are applied to an interface on the switch. You could apply IPv6-ACLs to VSAN interfaces, the management interface, Gigabit Ethernet interfaces on IPS modules and MPS-14/2 modules, and Ethernet PortChannel interfaces.

Tip Apply the IPv6-ACL on the interface closest to the source of the traffic.

When you are trying to block traffic from source to destination, you can apply an inbound IPv6-ACL to M0 on Switch 1 instead of an outbound filter to M1 on Switch 3 (see Figure 37-6).

Figure 37-6 Denying Traffic on the Inbound Interface

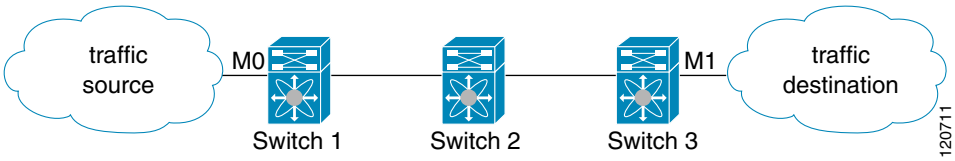

The **access-group** option controls access to an interface. Each interface can only be associated with one IPv6-ACL per direction. The ingress direction can have a different IPv6-ACL from the egress direction. The IPv6-ACL becomes active when applied to the interface.

Tip Create all conditions in an IPv6-ACL before applying it to the interface.

Caution If you apply an IPv6-ACL to an interface before creating it, all packets in that interface are dropped because the IPv6-ACL is empty.

The terms *in, out, source*, and *destination* are used as referenced by the switch.

- In—Traffic that arrives at the interface and goes through the switch; the source is where it transmitted from and the destination is where it is transmitted to (on the other side of the router).

Tip The IPv6-ACL applied to the interface for the ingress traffic affects both local and remote traffic.

- Out—Traffic that has already been through the switch and is leaving the interface; the source is where it transmitted from and the destination is where it is transmitted to.

Tip The IPv6-ACL applied to the interface for the egress traffic only affects local traffic.

To apply an IPv6-ACL to an interface using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **IP ACL** in the Physical Attributes pane.

You see the IP-ACL configuration in the Information pane. See Figure 37-7.

Figure 37-7 IP-ACL Configuration in the Physical Attributes Pane

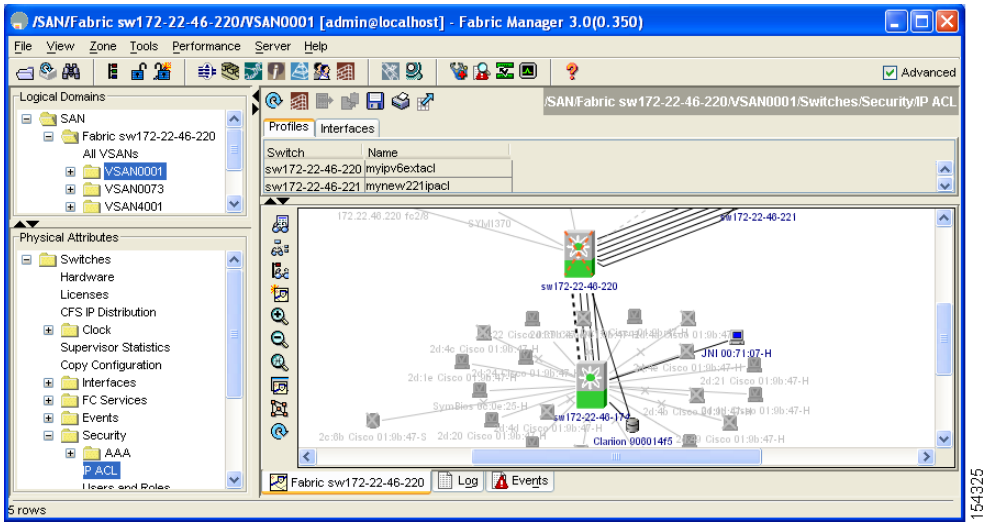

Step 2 Click the **Interfaces** tab.

You see a list of interfaces associated with the IP-ACLs. See Figure 37-8.

Figure 37-8 List of Interfaces Associated with the IP-ACLs

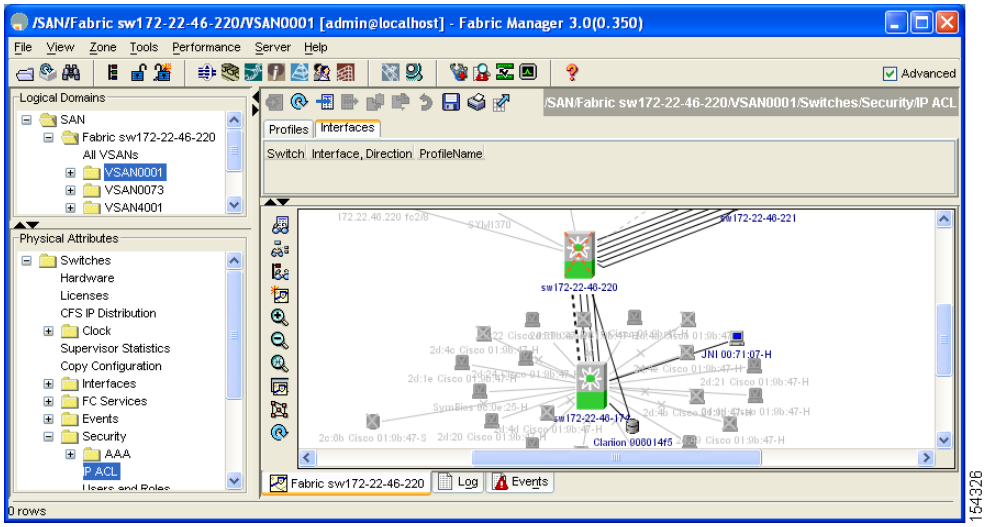

Step 3 Click **Create Row**.

You see the Create Interface dialog box.

Step 4 Optionally, select the switches you want to include in the IP-ACL by checking the check boxes next to the switch address in Fabric Manager.

Step 5 Set the interface you want associated with an IP-ACL in the Interface field.

Step 6 Choose the appropriate ProfileDirection radio button (**inbound**, **outbound**, **inboundIPv6**, or **outboundIPv6**).

Step 7 Enter the IPv6-ACL name in the Profile Name field.

Note This IPv6-ACL name must already have been created using the Create Profiles dialog box. If not, no filters will be enabled until you to go to the Create Profiles dialog box and create the profile.

Step 8 Click **Create** to associate the IPv6-ACL, or click **Close** to close the Create Interfaces dialog box without associating an access list.

You see the newly associated access list in the list of IPv6-ACLs.

CHAPTER 38

Configuring Certificate Authorities and Digital Certificates

Public Key Infrastructure (PKI) support provides the means for the Cisco MDS 9000 Family switches to obtain and use digital certificates for secure communication in the network. PKI support provides manageability and scalability for IPsec/IKE and SSH.

This chapter includes the following sections:

About CAs and Digital Certificates

This section provides information about certificate authorities (CAs) and digital certificates, and includes the following topics:

Purpose of CAs and Digital Certificates

CAs manage certificate requests and issue certificates to participating entities such as hosts, network devices, or users. The CAs provide centralized key management for the participating entities.

Digital signatures, based on public key cryptography, digitally authenticate devices and individual users. In public key cryptography, such as the RSA encryption system, each device or user has a key-pair containing both a private key and a public key. The private key is kept secret and is known only to the owning device or user only. However, the public key is known to everybody. The keys act as complements. Anything encrypted with one of the keys can be decrypted with the other. A signature is formed when data is encrypted with a sender's private key. The receiver verifies the signature by decrypting the message with the sender's public key. This process relies on the receiver having a copy of the sender's public key and knowing with a high degree of certainty that it really does belong to the sender and not to someone pretending to be the sender.

Digital certificates link the digital signature to the sender. A digital certificate contains information to identify a user or device, such as the name, serial number, company, department, or IP address. It also contains a copy of the entity's public key. The certificate is itself signed by a CA, a third party that is explicitly trusted by the receiver to validate identities and to create digital certificates.

To validate the signature of the CA, the receiver must first know the CA's public key. Normally this process is handled out-of-band or through an operation done at installation. For instance, most web browsers are configured with the public keys of several CAs by default. The Internet Key Exchange (IKE), an essential component of IPsec, can use digital signatures to scalably authenticate peer devices before setting up security associations.

Trust Model, Trust Points, and Identity CAs

The trust model used in PKI support is hierarchical with multiple configurable trusted CAs. Each participating entity is configured with a list of CAs to be trusted so that the peer's certificate obtained during the security protocol exchanges can be verified, provided it has been issued by one of the locally trusted CAs. To accomplish this, CA's self signed root certificate (or certificate chain for a subordinate CA) is locally stored. The process of securely obtaining a trusted CA's root certificate (or the entire chain in the case of a subordinate CA) and storing it locally is called *CA authentication* and is a mandatory step in trusting a CA.

The information about a trusted CA that is locally configured is called the *trust point* and the CA itself is called a *trust point CA*. This information consists of CA certificate (or certificate chain in case of a subordinate CA) and the certificate revocation checking information.

The MDS switch can also enroll with a trust point to obtain an identity certificate (for example, for IPsec/IKE). This trust point is called an *identity CA*.

RSA Key-Pairs and Identity Certificates

You can generate one or more RSA key-pairs and associate each RSA key-pair with a trust point CA where the MDS switch intends to enroll to obtain an identity certificate. The MDS switch needs only one identity per CA, which consists of one key-pair and one identity certificate per CA.

Cisco MDS SAN-OS allows you to generate RSA key-pairs with a configurable key size (or modulus). The default key size is 512. You can also configure an RSA key-pair label. The default key label is the switch fully qualified domain name (FQDN).

The following list summarizes the relationship between trust points, RSA key-pairs, and identity certificates:

- A trust point corresponds to a specific CA that the MDS switch trusts for peer certificate verification for any application (such as IKE or SSH).
- An MDS switch can have many trust points and all applications on the switch can trust a peer certificate issued by any of the trust point CAs.
- A trust point is not restricted to a specific application.
- An MDS switch enrolls with the CA corresponding to the trust point to obtain an identity certificate. You can enroll your switch with multiple trust points thereby obtaining a separate identity certificate from each trust point. The identity certificates are used by applications depending upon the purposes specified in the certificate by the issuing CA. The purpose of a certificate is stored in the certificate as certificate extensions.
- When enrolling with a trust point, you must specify an RSA key-pair to be certified. This key-pair must be generated and associated to the trust point before generating the enrollment request. The association between the trust point, key-pair, and identity certificate is valid until it is explicitly removed by deleting the certificate, key-pair, or trust point.
- The subject name in the identity certificate is the fully qualified domain name for the MDS switch.
- You can generate one or more RSA key-pairs on a switch and each can be associated to one or more trust points. But no more than one key-pair can be associated to a trust point, which means only one identity certificate is allowed from a CA.
- If multiple identity certificates (each from a distinct CA) have been obtained, the certificate that an application selects to use in a security protocol exchange with a peer is application specific (see the "IPsec Digital Certificate Support" section on page 39-14 and the "SSH Authentication Using Digital Certificates" section on page 33-18).
- You do not need to designate one or more trust points for an application. Any application can use any certificate issued by any trust point as long as the certificate purpose satisfies the application requirements.
- You do not need more than one identity certificate from a trust point or more than one key-pair to be associated to a trust point. A CA certifies a given identity (name) only once and does not issue multiple certificates with the same subject name. If you need more than one identity certificate for a CA, then define another trust point for the same CA, associate another key-pair to it, and have it certified, provided CA allows multiple certificates with the same subject name.

Multiple Trusted CA Support

An MDS switch can be configured to trust multiple CAs by configuring multiple trust points and associating each with a distinct CA. With multiple trusted CAs, you do not have to enroll a switch with the specific CA that issued a certificate to a peer. Instead, you configure the switch with multiple trusted CAs that it trusts. A switch can then use a configured trusted CA to verify certificates offered by a peer that were not issued by the same CA defined in the identity of the switch.

Configuring multiple trusted CAs allows two or more switches enrolled under different domains (different CAs) to verify the identity of each other when using IKE to set up IPsec tunnels.

PKI Enrollment Support

Enrollment is the process of obtaining an identity certificate for the switch that is used for applications like IPsec/IKE or SSH. It occurs between the switch requesting the certificate and the certificate authority.

The PKI enrollment process for a switch involves the following steps:

1. Generate an RSA private and public key-pair on the switch.
2. Generate a certificate request in standard format and forward it to the CA.
3. Manual intervention at the CA server by the CA administrator may be required to approve the enrollment request, when it is received by the CA.
4. Receive the issued certificate back from the CA, signed with the CA's private key.
5. Write the certificate into a nonvolatile storage area on the switch (bootflash).

Manual Enrollment Using Cut-and-Paste Method

Cisco MDS SAN-OS supports certificate retrieval and enrollment using a manual cut-and-paste method. Cut-and-paste enrollment literally means you must cut and paste the certificate requests and resulting certificates between the switch and the CA, as follows:

1. Create an enrollment certificate request, which is displayed in base64-encoded text form.
2. Cut and paste the encoded certificate request text in an e-mail message or in a web form and send it to the CA.
3. Receive the issued certificate (in base64-encoded text form) from the CA in an e-mail message or in a web browser download.
4. Cut and paste the issued certificate to the switch using the certificate import facility.

Note Fabric Manager does not support cut and paste. Instead, it allows the enrollment request (certificate signing request) to be saved in a file to be sent manually to the CA.

Multiple RSA Key-Pair and Identity CA Support

Multiple identity CA support enables the switch to enroll with more than one trust point. This results in multiple identity certificates; each from a distinct CA. This allows the switch to participate in IPsec and other applications with many peers using certificates issued by appropriate CAs that are acceptable to those peers.

The multiple RSA key-pair support feature allows the switch to maintain a distinct key pair for each CA with which it is enrolled. Thus, it can match policy requirements for each CA without conflicting the requirements specified by the other CAs, such as key length. The switch can generate multiple RSA key-pairs and associate each key pair with a distinct trust point. Thereafter, when the enrolling with a trust point, the associated key pair is used to construct the certificate request.

Peer Certificate Verification

The PKI support on an MDS switch provides the means to verify peer certificates. The switch verifies certificates presented by peers during security exchanges pertaining to applications, such as IPsec/IKE and SSH. The applications verify the validity of the peer certificates presented to them. The peer certificate verification process involves the following steps:

- Verifies that the peer certificate is issued by one of the locally trusted CAs.
- Verifies that the peer certificate is valid (not expired) with respect to current time.
- Verifies that the peer certificate is not yet revoked by the issuing CA.

For revocation checking, two methods are supported: certificate revocation list (CRL) and Online Certificate Status Protocol (OCSP). A trust point uses one or both of these methods to verify that the peer certificate has not been revoked.

CRL Downloading, Caching, and Checking Support

Certificate revocation lists (CRLs) are maintained by CAs to give information of prematurely revoked certificates, and the CRLs are published in a repository. The download URL is made public and also specified in all issued certificates. A client verifying a peer's certificate should obtain the latest CRL from the issuing CA and use it to determine if the certificate has been revoked. A client can cache the CRLs of some or all of its trusted CAs locally and use them later if necessary until the CRLs expire.

Cisco MDS SAN-OS allows the manual configuration of pre-downloaded of CRLs for the trust points, and then caches them in the switch bootflash (cert-store). During the verification of a peer certificate by IPsec or SSH, the issuing CA's CRL is consulted only if the CRL has already been cached locally and the revocation checking is configured to use CRL. Otherwise, CRL checking is not performed and the certificate is considered to be not revoked if no other revocation checking methods are configured. This mode of CRL checking is called CRL optional.

OCSP Support

Online Certificate Status Protocol (OCSP) facilitates online certificate revocation checking. You can specify an OCSP URL for each trust point. Applications choose the revocation checking mechanisms in a specified order. The choices are CRL, OCSP, none, or a combination of these methods.

Import and Export Support for Certificates and Associated Key Pairs

As part of the CA authentication and enrollment process, the subordinate CA certificate (or certificate chain) and identity certificates can be imported in standard PEM (base64) format.

The complete identity information in a trust point can be exported to a file in the password-protected PKCS#12 standard format. It can be later imported to the same switch (for example, after a system crash) or to a replacement switch. The information in a PKCS#12 file consists of the RSA key-pair, the identity certificate, and the certificate chain.

Configuring CAs and Digital Certificates

This section describes the tasks you must perform to allow CAs and digital certificates your Cisco MDS switch device to interoperate. This section includes the following topics:

Configuring the Host Name and IP Domain Name

You must configure the host name and IP domain name of the switch if they are not already configured. This is required because switch FQDN is used as the subject in the identity certificate. Also, the switch FQDN is used as a default key label when none is specified during key-pair generation. For example, a certificate named SwitchA.example.com is based on a switch host name of SwitchA and a switch IP domain name of example.com.

Caution Changing the host name or IP domain name after generating the certificate can invalidate the certificate.

To configure the host name and IP domain name, refer to the *Cisco MDS 9000 SAN-OS CLI Configuration Guide*.

Creating an RSA Key-Pair

RSA key-pairs are used to sign and/or encrypt and decrypt the security payload during security protocol exchanges for applications such as IKE/IPsec and SSH, and they are required before you can obtain a certificate for your switch.

To generate an RSA key-pair using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **PKI** in the Information pane.

Step 2 Click the **RSA Key-Pair** tab.

You see the information in Figure 38-1.

Figure 38-1 PKI RSA Key-Pair Information

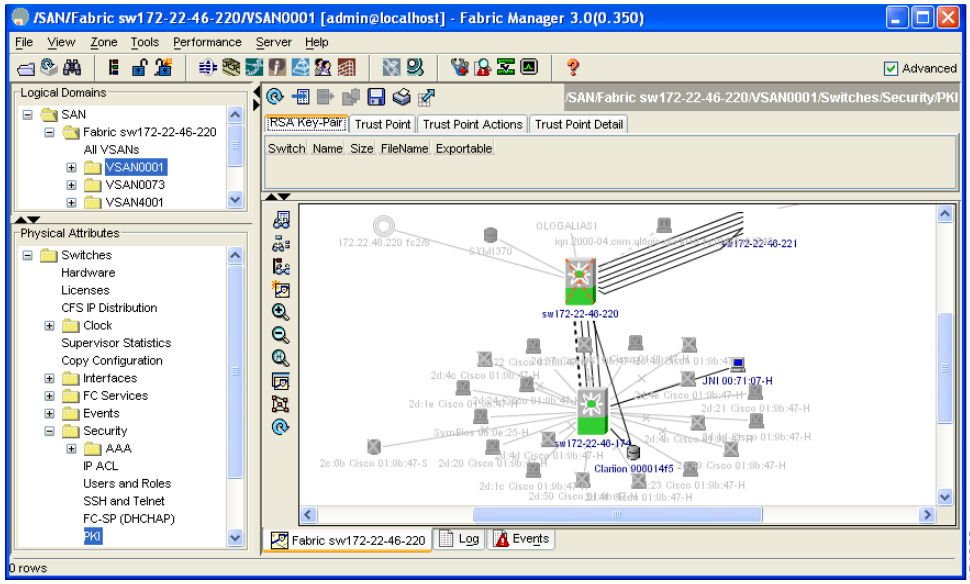

Step 3 Click **Create Row**.

You see the **Create RSA Key-Pair** dialog shown in Figure 38-2.

Figure 38-2 Create RSA Key-Pair Dialog Box

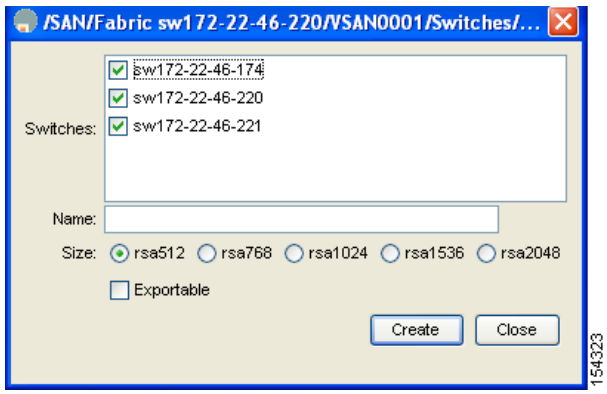

Step 4 Select the switches for which you want to create the RSA key-pair.

Step 5 Assign a name to the RSA key-pair.

Step 6 Select the Size or modulus values. Valid modulus values are 512, 768, 1024, 1536, and 2048.

Note The security policy (or requirement) at the local site (MDS switch) and at the CA (where enrollment is planned) are considered in deciding the appropriate key modulus.

Note The maximum number of key-pairs you can configure on a switch is 16.

Step 7 Check the **Exportable** check box if you want the key to be exportable.

Caution The exportability of a key-pair cannot be changed after key-pair generation.

Note Only exportable key-pairs can be exported in PKCS#12 format.

Step 8 Click **Create** to create the RSA Key-Pair or click **Close** if you do not want to create the RSA key-pair.

Creating a Trust Point CA

To create a trust point CA using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point** tab in the Information Pane.

You see the information shown in Figure 38-3.

Figure 38-3 Trust Point Tab

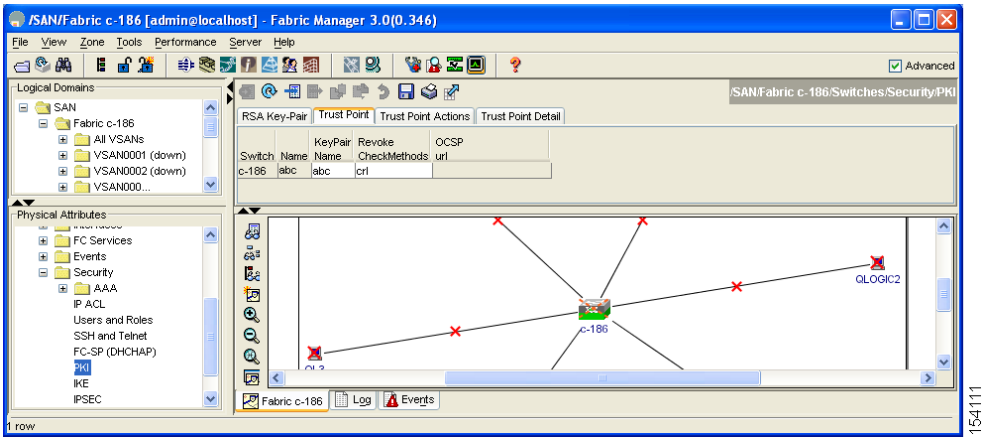

Step 3 Click **Create Row**.

You see the **Create Trust Point** dialog box in Figure 38-4.

Figure 38-4 Create Trust Point Dialog Box

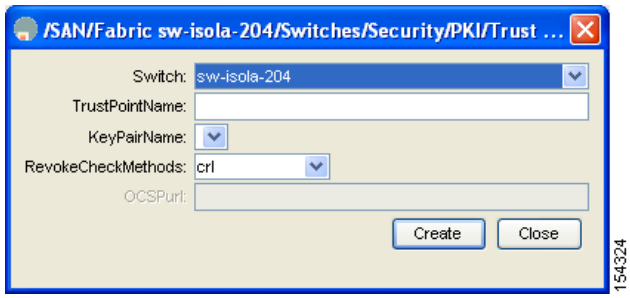

Step 4 Select the switch for which you are creating the trust point CA from the **Switch** drop-down menu.

Step 5 Assign a name to the trust point CA.

Step 6 Select a key-pair name to be associated with this trust point for enrollment. It was generated earlier in the "Creating an RSA Key-Pair" section on page 38-7. Only one RSA key-pair can be specified per CA.

Step 7 From the RevokeCheckMethod drop-down menu, select the certificate revocation method that you would like to use (see Figure 38-4). You can use CRL, OCSP, CRL OCSP, or OCSP CRL to check for certificate revocation. The CRL OCSP option checks for revoked certificates first in the locally stored CRL. If not found, the switch uses OCSP to check the revoked certificates on the URL specified in Step 7.

Step 8 Enter the OCSP URL if you selected an OCSP certificate revocation method.

Note The OSCP URL must be configured before configuring the revocation checking method.

Step 9 Click **Create** to successfully create the trust point CA or click **Close** to close the Create Trust Point dialog without creating the trust point CA.

Copying Files to Bootflash

To copy files to bootflash using Device Manager, follow these steps:

Step 1 Click **Admin > Flash Files**.

Step 2 Select bootflash in the Device field.

You see a list of flash files in the dialog box shown in Figure 38-5.

Figure 38-5 Flash Files

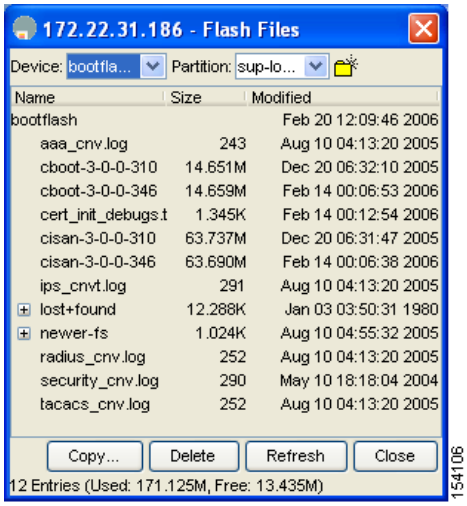

Step 3 Click **Copy**.

You see the Copy Files dialog box shown in Figure 38-6.

Figure 38-6 Copy Files Dialog Box

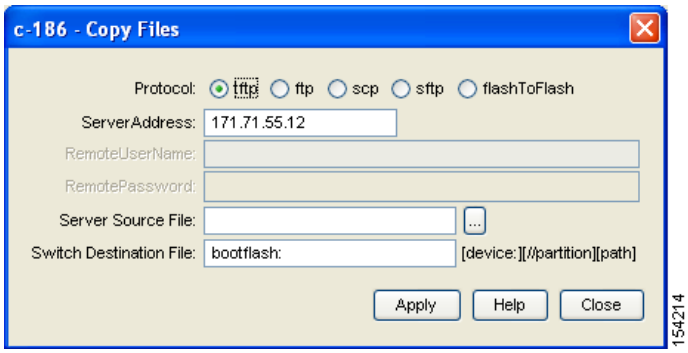

Step 4 Select **tftp** as the Protocol field.

Step 5 Click the Browse button (**...**) to locate the appropriate file to copy to bootflash.

Step 6 Click **Apply** to apply these changes or click **Close** if you do not want to proceed with these changes.

Authenticating the CA

The configuration process of trusting a CA is complete only when the CA is authenticated to the MDS switch. The switch must authenticate the CA. It does this by obtaining the self-signed certificate of the CA in PEM format, which contains the public key of the CA. Because the certificate of the CA is self-signed (the CA signs its own certificate) the public key of the CA should be manually authenticated by contacting the CA administrator to compare the fingerprint of the CA certificate.

Note If the CA being authenticated is not a self-signed CA (that is, it is a subordinate CA to another CA, which itself may be a subordinate to yet another CA, and so on, finally ending in a self-signed CA), then the full list of the CA certificates of all the CAs in the certification chain needs to be input during the CA authentication step. This is called the *CA certificate chain* of the CA being authenticated. The maximum number of certificates in a CA certificate chain is 10.

To authenticate a CA using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information pane.

You see the information shown in Figure 38-7.

Figure 38-7 Trust Point Actions Tab

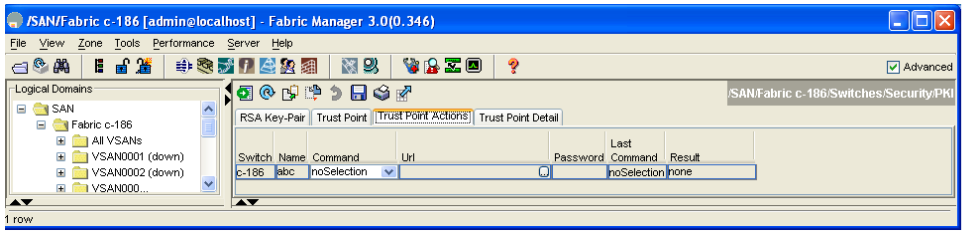

Step 3 Click the Command field drop-down menu and select the appropriate option. Available options are **caauth**, **cadelete**, **certreq**, **certimport**, **certdelete**, **pkcs12import**, and **pkcs12export**. The **caauth** option is provided to authenticate a CA and install its CA certificate or certificate chain in a trust point.

Step 4 Click the Browse (**...**) button in the URL field and select the appropriate import certificate file from the Bootflash Files dialog box. It is the file name containing the CA certificate or chain in the bootflash:filename format.

Note You can authenticate a maximum of 10 trust points to a specific CA.

Note If you do not see the required file in the Import Certificate dialog box, make sure that you copy the file to bootflash. See "Copying Files to Bootflash" section on page 10.

Step 5 Click **Apply Changes** to save the changes.

Authentication is then confirmed or not confirmed depending on whether or not the certificate can be accepted after manual verification of its fingerprint.

Note For subordinate CA authentication, the full chain of CA certificates ending in a self-signed CA is required because the CA chain is needed for certificate verification as well as for PKCS#12 format export.

Confirming CA Authentication

As mentioned in step 5 of "Authenticating the CA" section on page 38-11, CA authentication is required to be followed by CA confirmation in order to accept the CA certificate based on its fingerprint verification.

To confirm CA authentication using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information Pane.

Step 3 Make a note of the CA certificate fingerprint displayed in the IssuerCert FingerPrint column for the trust point row in question. Compare the CA certificate fingerprint with the fingerprint already communicated by the CA (obtained from the CA web site).

If the fingerprints match exactly, accept the CA with the **certconfirm** command in the Command drop-down menu. Otherwise, reject the CA with the **certnoconfirm** command.

Step 4 If you selected **certconfirm** in step 3, click Command and select the **certconfirm** action from the drop-down menu. Click **Apply Changes.**

If you selected **certnoconfirm** in step 3, click Command and select the **certnoconfirm** action drop-down menu. Click **Apply Changes**.

Configuring Certificate Revocation Checking Methods

During security exchanges with a client (for example, an IKE peer or SSH user), the MDS switch performs the certificate verification of the peer certificate sent by the client and the verification process may involve certificate revocation status checking.

You can use different methods for checking for revoked sender certificates. You can configure the switch to check the CRL downloaded from the CA (see the "Configuring a CRL" section on page 38-17), you can use OSCP if it is supported in your network, or both. Downloading the CRL and checking locally does not generate traffic in your network. However, certificates can be revoked between downloads and your switch would not be aware of the revocation. OCSP provides the means to check the current CRL on the CA. However, OCSP can generate network traffic that can impact network efficiency. Using both local CRL checking and OCSP provides the most secure method for checking for revoked certificates.

Note You must authenticate the CA before configuring certificate revocation checking.

Fabric Manager allows you to configure certificate revocation checking methods when you are creating a trust point CA. See "Creating a Trust Point CA" section on page 38-8.

Generating Certificate Requests

You must generate a request to obtain identity certificates from the associated trust point CA for each of your switch's RSA key-pairs. You must then cut and paste the displayed request into an e-mail message or in a website form for the CA.

To generate a request for signed certificates from the CA using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information pane in Figure 38-8.

Figure 38-8 Trust Point Actions Tab

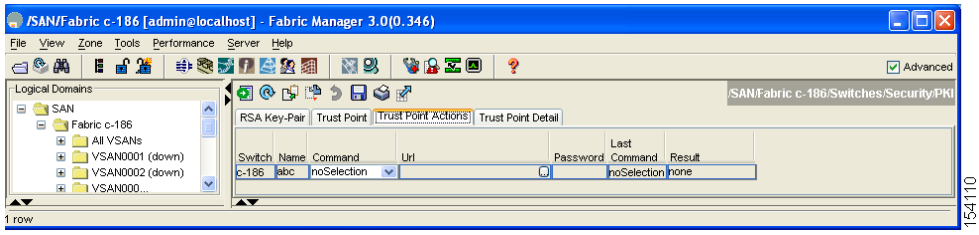

Step 3 Select the **certreq** option from the Command drop-down menu. This generates a pkcs#10 certificate signing request (CSR) needed for an identity certificate from the CA corresponding to this trust point entry. This entry requires an associated key-pair. The CA certificate or certificate chain should already be configured through the **caauth** action. See "Authenticating the CA" section on page 38-11.

Step 4 Enter the output file name for storing the generated certificate request. It will be used to store the CSR generated in PEM format. Use the format bootflash:filename. This CSR should be submitted to the CA to get the identity certificate. Once the identity certificate is obtained, it should be installed in this trust point. See "Installing Identity Certificates" section on page 38-14.

Step 5 Enter the *challenge* password to be included in the CSR.

Note The challenge password is not saved with the configuration. This password is required in the event that your certificate needs to be revoked, so you must remember this password.

Step 6 Click **Apply Changes** to save the changes.

Installing Identity Certificates

You receive the identity certificate from the CA by e-mail or through a web browser in base64-encoded text form. You must install the identity certificate from the CA by cutting and pasting the encoded text.

To install an identity certificate received from the CA using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab, in the Information pane.

Step 3 Select the **certimport** option from the Command drop-down menu to import an identity certificate in this trust point. The identity certificate is obtained from the corresponding CA for a CSR generated previously (see "Generating Certificate Requests" section on page 38-13).

Note The identity certificate should be available in PEM format in a file in bootflash.

Step 4 Enter the name of the certificate file that should have been copied to bootflash in the URL field in the bootflash:filename format.

Step 5 Click **Apply Changes** to save your changes.

If successful, the values of the identity certificate and its related objects, like the certificate file name, are automatically updated with the appropriate values as per the corresponding attributes in the identity certificate.

Saving Your Configuration

Save your work when you make configuration changes or the information is lost when you exit.

To save your configuration using Fabric Manager, follow these steps:

Step 1 Expand **Switches** and then select **Copy Configuration** in the Physical Attributes pane.

Step 2 Select the switch configuration including the RSA key-pairs and certificates.

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

Ensuring Trust Point Configurations Persist Across Reboots

The trust point configuration is a normal Cisco SAN-OS configuration that persists across system reboots only if you copy it explicitly to the startup configuration. The certificates, key-pairs, and CRL associated with a trust point are automatically persistent if you have already copied the trust point configuration in the startup configuration. Conversely, if the trust point configuration is not copied to the startup configuration, the certificates, key-pairs, and CRL associated with it are not persistent since they require the corresponding trust point configuration after a reboot. Always copy the running configuration to the startup configuration to ensure the that the configured certificates, key-pairs, and CRLs are persistent. Also, save the running configuration after deleting a certificate or key-pair to ensure the deletions permanent.

The certificates and CRL associated with a trust point automatically become persistent when imported (that is, without an explicitly copying to the startup configuration) if the specific trust point is already saved in startup configuration.

We also recommend that you create a password protected backup of the identity certificates nd save it to an external server (see the "Exporting and Importing Identity Information in PKCS#12 Format" section on page 38-15).

Note Copying the configuration to an external server does include the certificates and key-pairs.

Monitoring and Maintaining CA and Certificates Configuration

The tasks in the section are optional. This section includes the following topics:

Exporting and Importing Identity Information in PKCS#12 Format

You can export the identity certificate along with the RSA key-pair and CA certificate chain of a trust point to a PKCS#12 file for backup purposes. You can later import the certificate and RSA key-pair to recover from a system crash on your switch.

Note Only bootflash:*filename* format is supported to specify the export URL.

To export a certificate and key pair to a PKCS#12-formatted file using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information Pane.

Step 3 Select the **pkcs12export** option in the Command drop-down menu to export the key-pair, identity certificate, and the CA certificate or certificate chain in PKCS#12 format from the selected trust point.

Figure 38-9 Pkcs12export Option Exports a Key-Pair

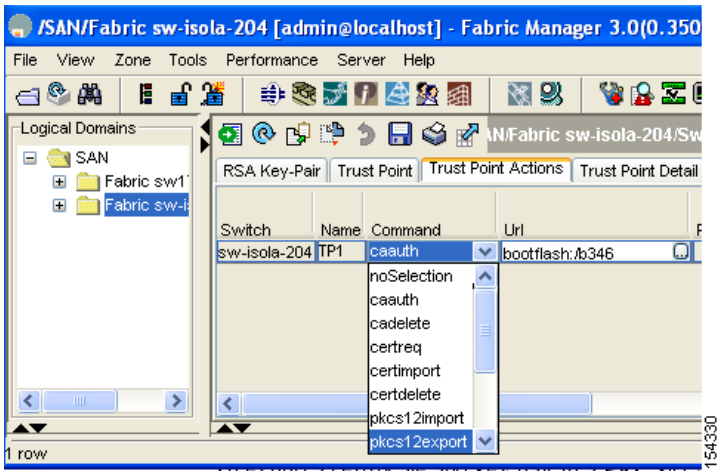

Step 4 Enter the output file name as bootflash:filename to store the exported PKCS#12 identity.

Step 5 Enter the required password. The password is set for encoding the PKCS#12 data. On successful completion, the exported data is available in bootflash in the specified file.

Step 6 Click **Apply Changes** to save the changes.

To import a certificate and key pair formatted as a PKCS#12 formatted file, follow these steps:

Step 1 Expand **Switches > Security** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information pane.

Step 3 Select the **pkcs12import** option from the Command drop-down menu to import the key-pair, identity certificate, and the CA certificate or certificate chain in the PKCS#12 format to the selected trust point.

Step 4 Enter the input in the bootflash:filename format, containing the PKCS#12 identity.

Step 5 Enter the required password. The password is set for decoding the PKCS#12 data. On completion, the imported data is available in bootflash in the specified file.

Step 6 Click **Apply Changes** to save the changes.

On completion the trust point is created in the RSA key-pair table corresponding to the imported key-pair. The certificate information is updated in the trust point.

Note The trust point should be empty (no RSA key-pair associated with it and no CA is associated with it using CA authentication) for the PKCS#12 import to succeed.

Configuring a CRL

To configure the CRL from a file to a trust point using Fabric Manager, follow these steps:

Step 1 Click **Switches > Security > PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information pane.

Step 3 Select the **crlimport** option from the Command drop-down menu to import the CRL to the selected trust point.

Step 4 Enter the input file name with the CRL in the bootflash:filename format, in the URL field.

Step 5 Click **Apply Changes** to save the changes.

Deleting Certificates from the CA Configuration

You can delete the identity certificates and CA certificates that are configured in a trust point. You must first delete the identity certificate, followed by the CA certificates. Then after deleting the identity certificate, you can disassociate the RSA key-pair from a trust point. The certificate deletion is necessary to remove expired or revoked certificates, certificates whose key-pairs are compromised (or suspected to be compromised) or CAs that are no longer trusted.

To delete the CA certificate (or the entire chain in the case of a subordinate CA) from a trust point using Fabric Manager, follow these steps:

Step 1 Click **Switches > Security > PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information pane.

Step 3 Select the **cadelete** option from the Command drop-down menu to delete the identity certificate from a trust point.

Note If the identity certificate being deleted is the last-most or only identity certificate in the device, you must use the **forcecertdelete** action to delete it. This ensures that the administrator does not mistakenly delete the last-most or only identity certificate and leave the applications (such as IKE and SSH) without a certificate to use.

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

To delete the identity certificate, click the **Trust Point Actions** tab and select the **certdelete** or **forcecertdelete** in the Command drop-down menu.

Deleting RSA Key-Pairs from Your Switch

Under certain circumstances you may want to delete your switch's RSA key-pairs. For example, if you believe the RSA key-pairs were compromised in some way and should no longer be used, you should delete the key pairs.

To delete RSA key-pairs from your switch, follow these steps:

Step 1 Expand **Switches** > **Security** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **RSA Key-Pair** tab in the Information pane.

Step 3 Click **Delete Row**.

Step 4 Click **Yes** or **No** in the Confirmation dialog box.

Note After you delete RSA key-pairs from a switch, ask the CA administrator to revoke your switch's certificates at the CA. You must supply the challenge password you created when you originally requested the certificates. See "Generating Certificate Requests" section on page 38-13.

Example Configurations

This section shows an example of the tasks you can use to configure certificates and CRLs on the Cisco MDS 9000 Family switches using the Microsoft Windows Certificate server.

Configuring Certificates on the MDS Switch

To configure certificates on an MDS switch using Fabric Manager, follow these steps:

Step 1 Choose **Switches** and set the LogicalName field to configure the switch host name.

Step 2 Choose **Switches > Interfaces > Management > DNS** and set the DefaultDomainName field to configure the DNS domain name for the switch.

Step 3 To create an RSA key-pair for the switch, follow these steps:

a. Choose **Switches > Security > PKI** and select the RSA **Key-Pair** tab.

b. Click **Create Row** and set the name and size field.

c. Check the **Exportable** check box and click **Create**.

Step 4 To create a trust point and associate the RSA key-pairs with it, follow these steps:

a. Choose S**witches > Security > PKI** and select the **Trustpoints** tab.

b. Click **Create Row** and set the TrustPointName field.

c. Select the RSA key-pairs from the KeyPairName drop-down menu.

d. Select the certificates revocation method from the CARevoke drop-down menu.

e. Click **Create**.

Step 5 Choose **Switches > Copy Configuration** and click **Apply Changes** to copy the running to startup configuration and save the trustpoint and key pair.

Step 6 Download the CA certificate from the CA that you want to add as the trustpoint CA.

Step 7 To authenticate the CA that you want to enroll to the trust point, follow these steps:

a. Using Device Manager, choose **Admin > Flash Files** and select **Copy** and tftp copy the CA certificate to bootflash.

b. Using Fabric Manager, choose **Switches > Security > PKI** and select the **TrustPoint Actions** tab.

c. Select **cauth** from the Command drop-down menu.

d. Click **...** in the URL field and select the CA certificate from bootflash.

e. Click **Apply Changes** to authenticate the CA that you want to enroll to the trust point.

f. Click the **Trust Point Actions** tab in the Information Pane.

g. Make a note of the CA certificate fingerprint displayed in the IssuerCert FingerPrint column for the trust point row in question. Compare the CA certificate fingerprint with the fingerprint already communicated by the CA (obtained from the CA web site). If the fingerprints match exactly, accept the CA by performing the **certconfirm** trust point action. Otherwise, reject the CA by performing the **certnoconfirm** trust point action.

h. If you select **certconfirm** in step g, select the **Trust Point Actions** tab, select **certconfirm** from the command drop-down menu and then click **Apply Changes**.

i. If you select **certnoconfirm** in step g, If you select **certconfirm** in step g, select the **Trust Point Actions** tab, select the **certnoconfirm** from the command drop-down menu and then click **Apply Changes**.

Step 8 To generate a certificate request for enrolling with that trust point, follow these steps:

a. Select the **Trust Point Actions** tab in the Information pane.

b. Select **certreq** from the Command drop-down menu. This generates a pkcs#10 certificate signing request (CSR) needed for an identity certificate from the CA corresponding to this trust point entry.

c. Enter the output file name for storing the generated certificate request. It should be specified in the bootflash:filename format and will be used to store the CSR generated in PEM format.

d. Enter the *challenge* password to be included in the CSR. The challenge password is not saved with the configuration. This password is required in the event that your certificate needs to be revoked, so you must remember this password.

e. Click **Apply Changes** to save the changes.

Step 9 Request an identity certificate from the CA.

Note The CA may require manual verification before issuing the identity certificate.

Step 10 To import the identity certificate, follow these steps:

a. Using Device Manager, choose **Admin > Flash Files** and select **Copy** and tftp copy the CA certificate to bootflash.

b. Using Fabric Manager, choose **Switches > Security > PKI** and select the **TrustPoint Actions** tab.

c. Select the **certimport** option from the Command drop-down menu to import an identity certificate in this trust point.

Note The identity certificate should be available in PEM format in a file in bootflash.

d. Enter the name of the certificate file which was copied to bootflash, in the URL field in the bootflash:filename format.

e. Click **Apply Changes** to save your changes.

If successful, the values of the identity certificate and its related objects, like the certificate file name, are automatically updated with the appropriate values as per the corresponding attributes in the identity certificate.

Downloading a CA Certificate

To download a CA certificate from the Microsoft Certificate Services web interface, follow these steps:

Step 1 Select the **Retrieve the CA certificate or certificate revocation task** radio button in the Microsoft Certificate Services web interface and click the **Next** button.

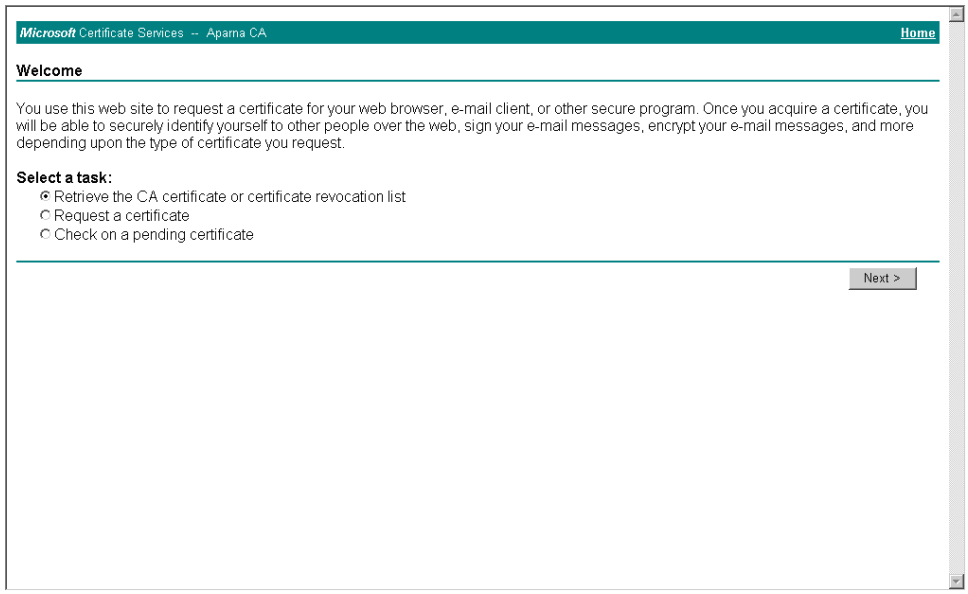

Step 2 Select the CA certificate file to download from the displayed list. Click the **Base 64 encoded** radio button, and click the **Download CA certificate** link.

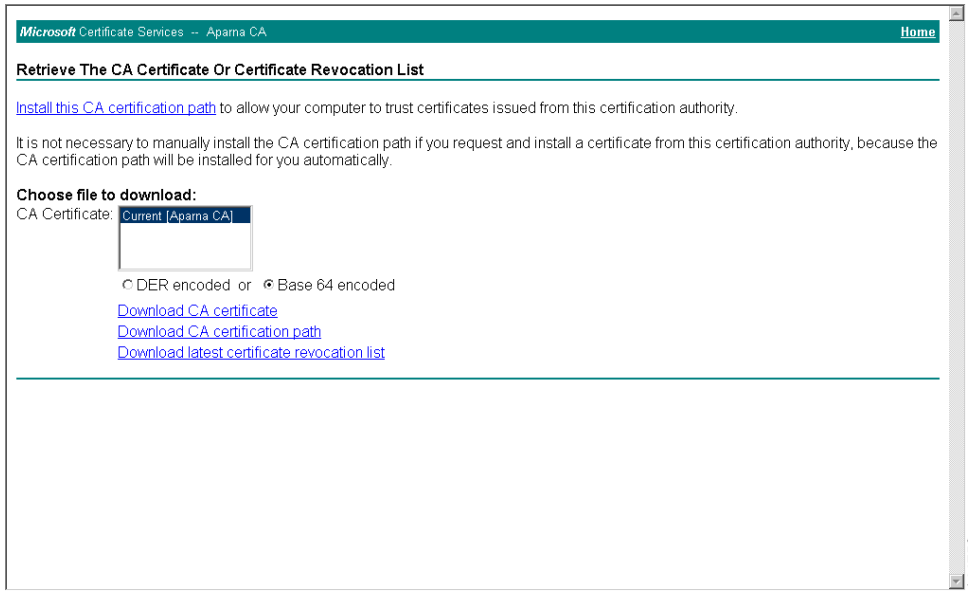

Step 3 Click the **Open** button in the File Download dialog box.

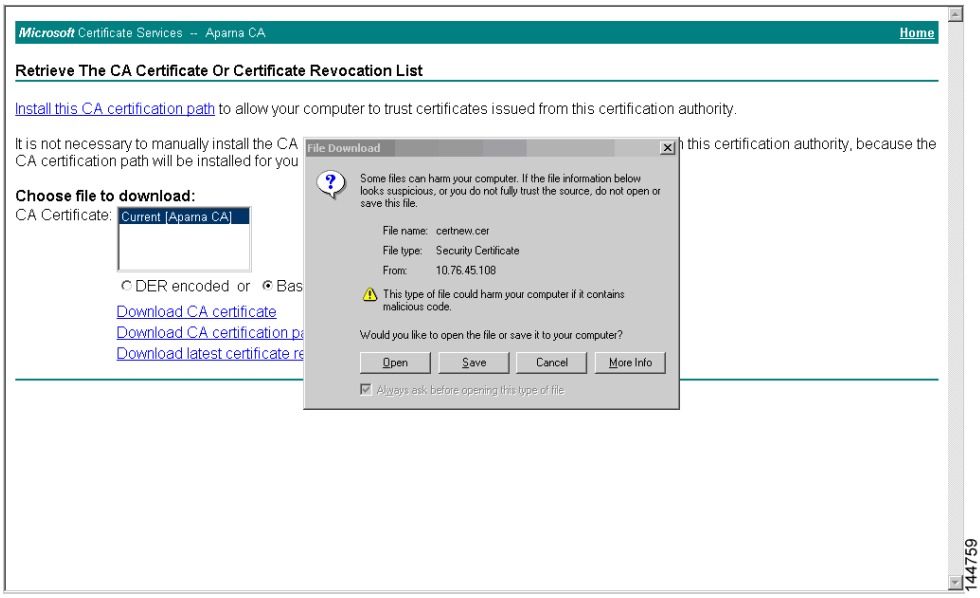

Step 4 Click the **Copy to File** button in the Certificate dialog box and click **OK**.

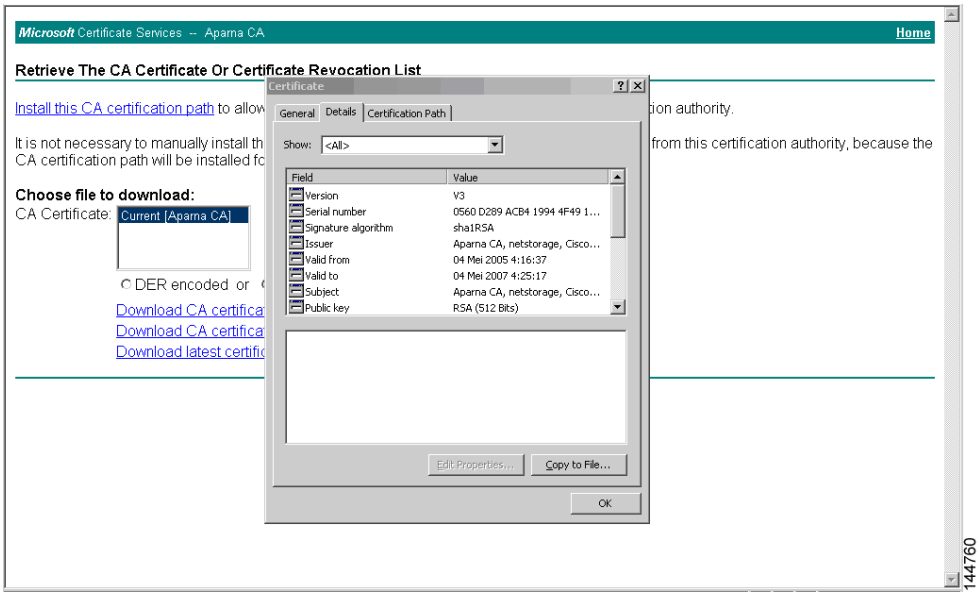

Step 5 Select the **Base-64 encoded X.509 (CER)** on the Certificate Export Wizard dialog box and click **Next**.

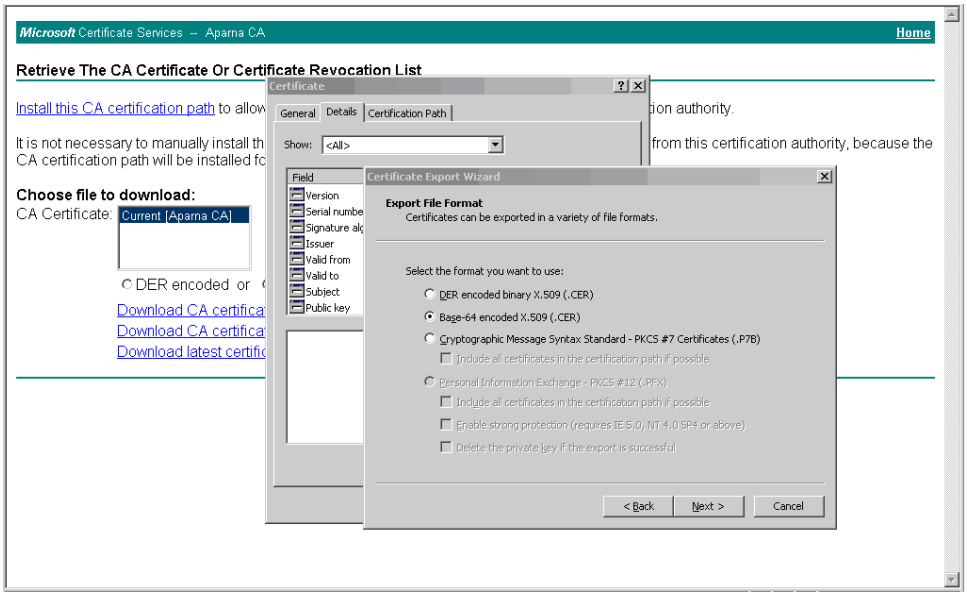

Step 6 Enter the destination file name in the File name: text box on the Certificate Export Wizard dialog box and click **Next**.

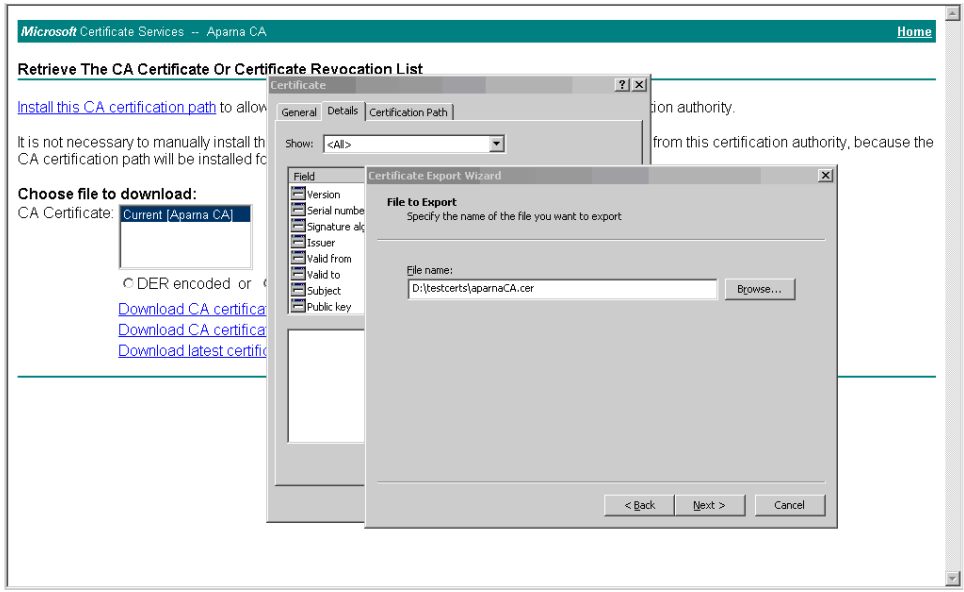

Step 7 Click the **Finish** button on the Certificate Export Wizard dialog box.

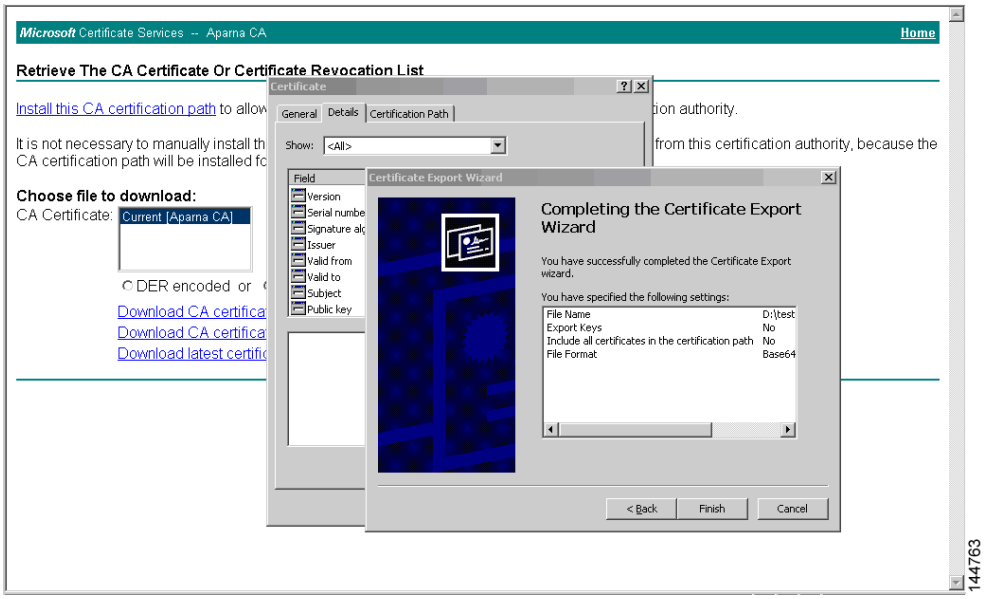

Step 8 Display the CA certificate stored in Base-64 (PEM) format.

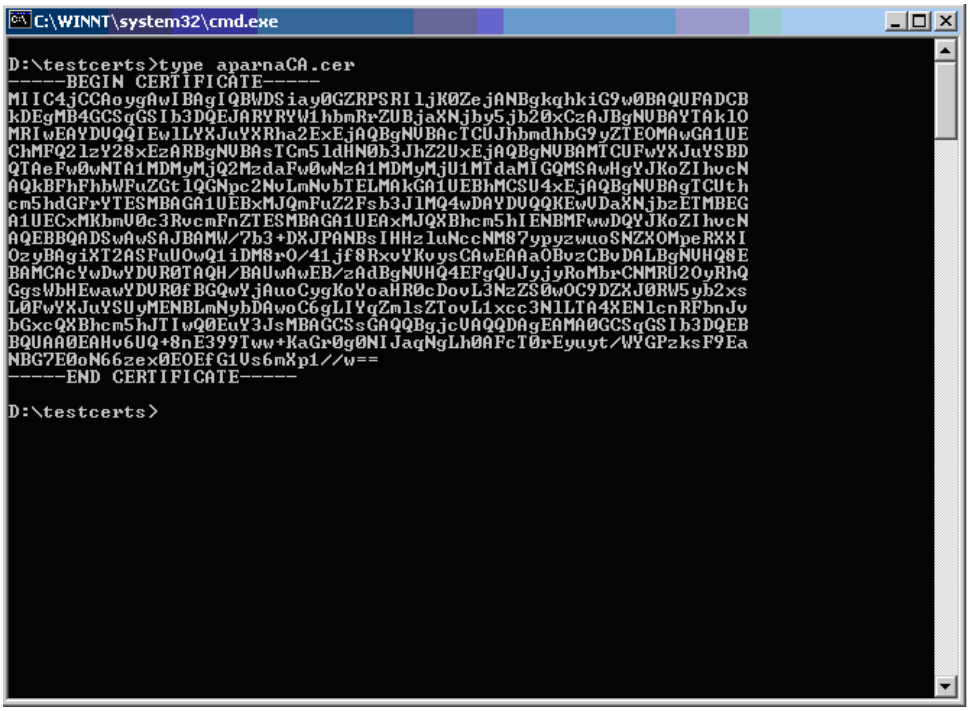

Requesting an Identity Certificate

To request an identify certificate from a Microsoft Certificate server using a PKCS#10 certificate signing request (CRS), follow these steps:

Step 1 Select the **Request a certificate** radio button on the Microsoft Certificate Services web interface and click **Next**.

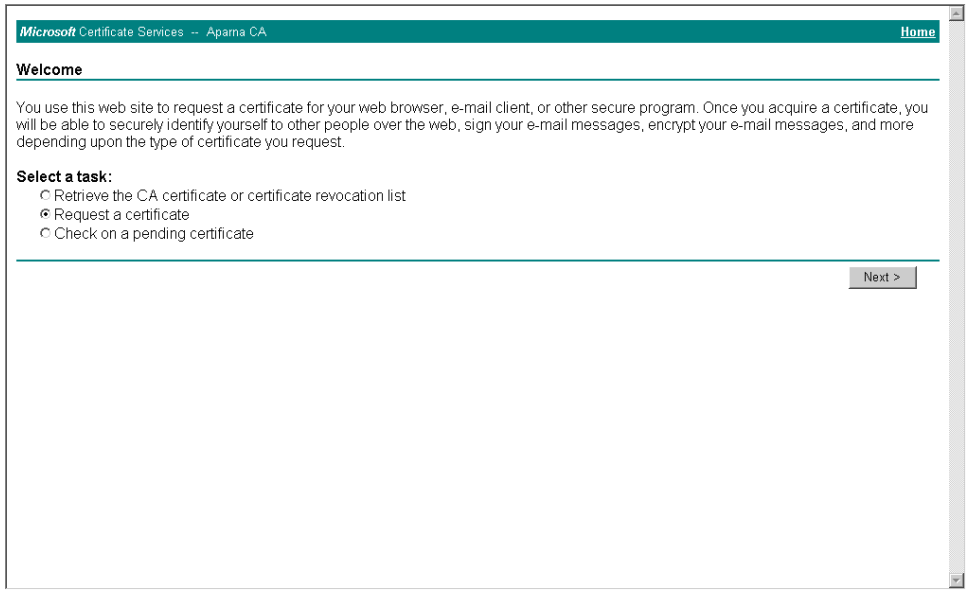

Step 2 Select the **Advanced Request** radio button and click **Next**.

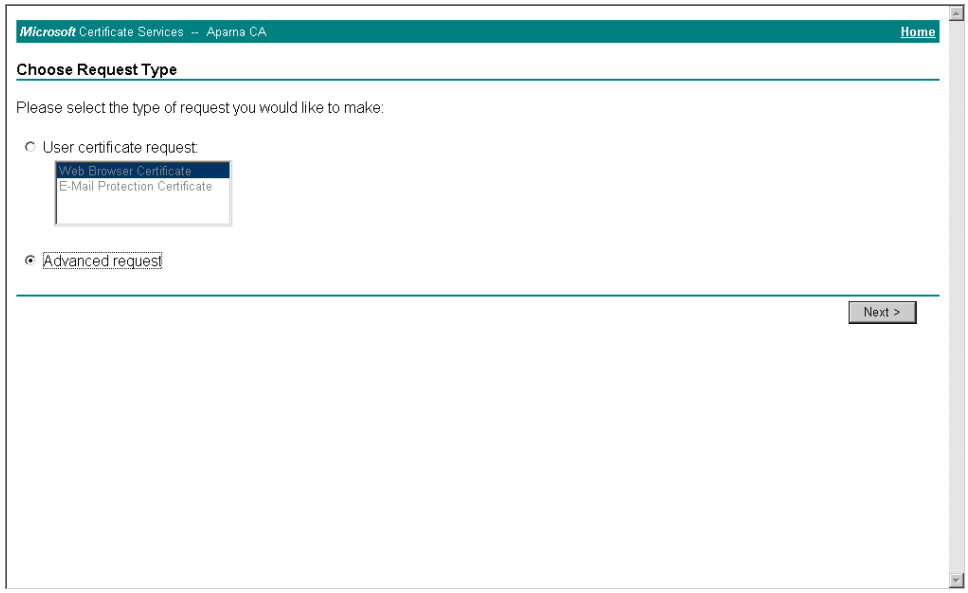

Step 3 Select the **Submit a certificate request using a base64 encoded PKCS#10 file or a renewal request using a base64 encoded PKCS#7 file** radio button and click **Next**.

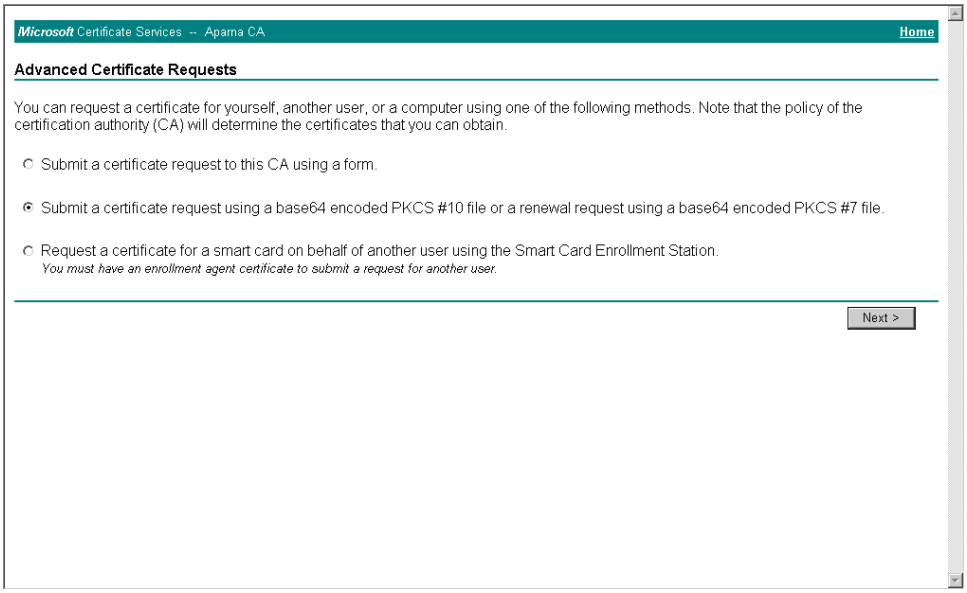

Step 4 Paste the base64 PKCS#10 certificate request in the Saved Request text box and click **Next**. The certificate request is copied from the MDS switch console (see the "Generating Certificate Requests" section on page 38-13 and "Downloading a CA Certificate" section on page 38-21)

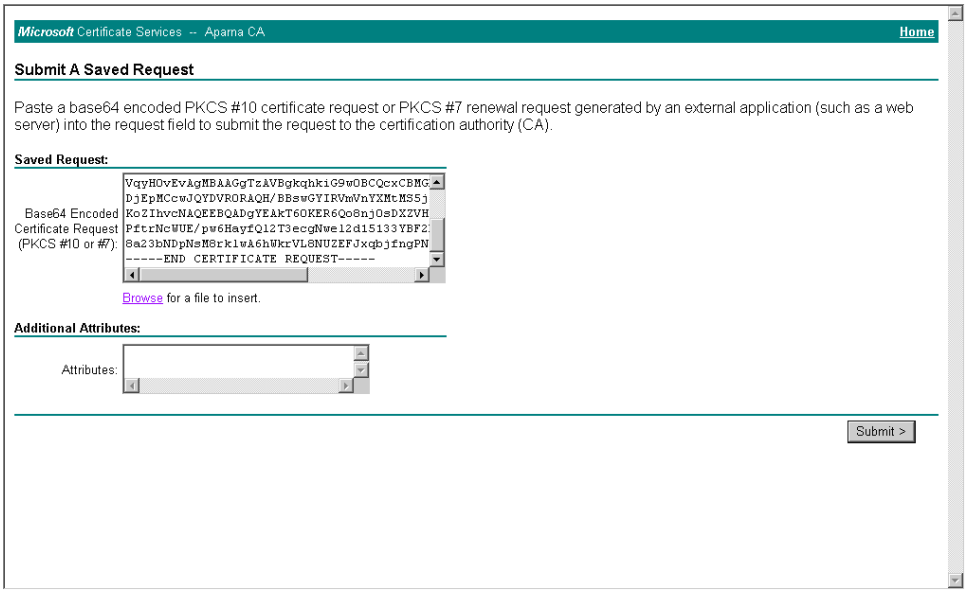

Step 5 Wait one or two days until the certificate is issued by the CA administrator.

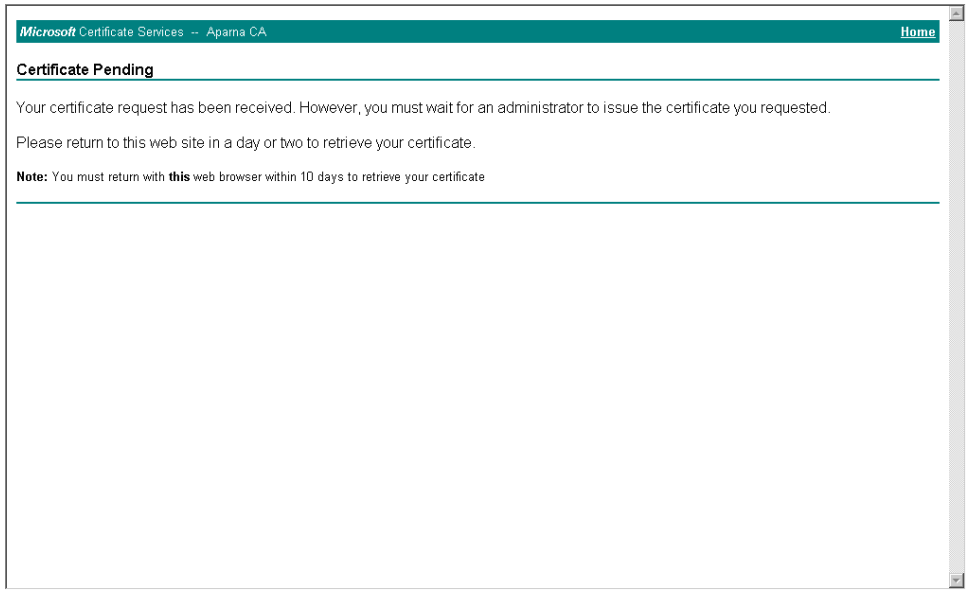

Step 6 The CA administrator approves the certificate request.

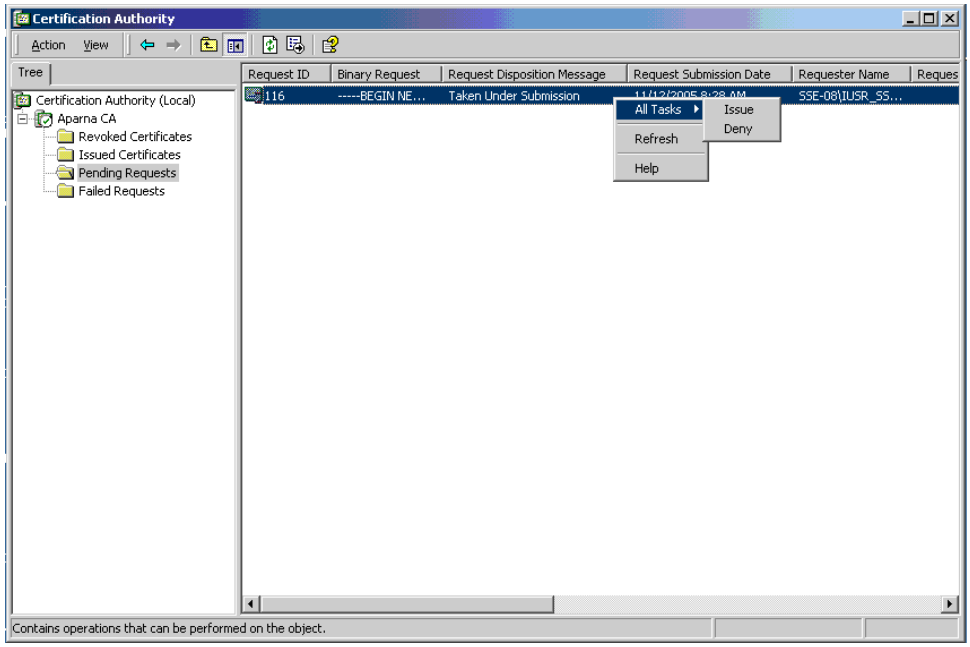

Step 7 Select the **Check on a pending certificate** radio button on the Microsoft Certificate Services web interface and click **Next**.

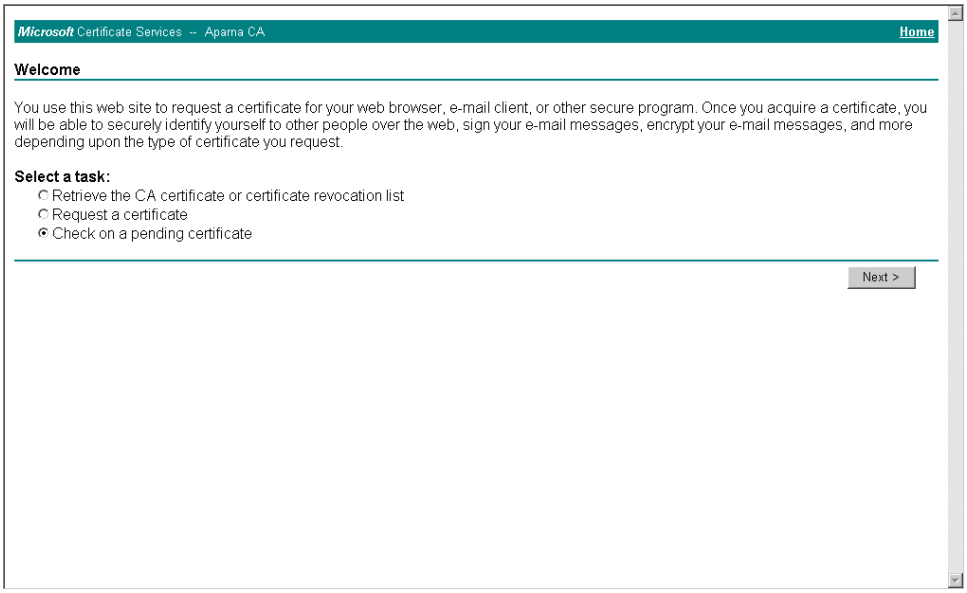

Step 8 Select the certificate request you want to check and click **Next**.

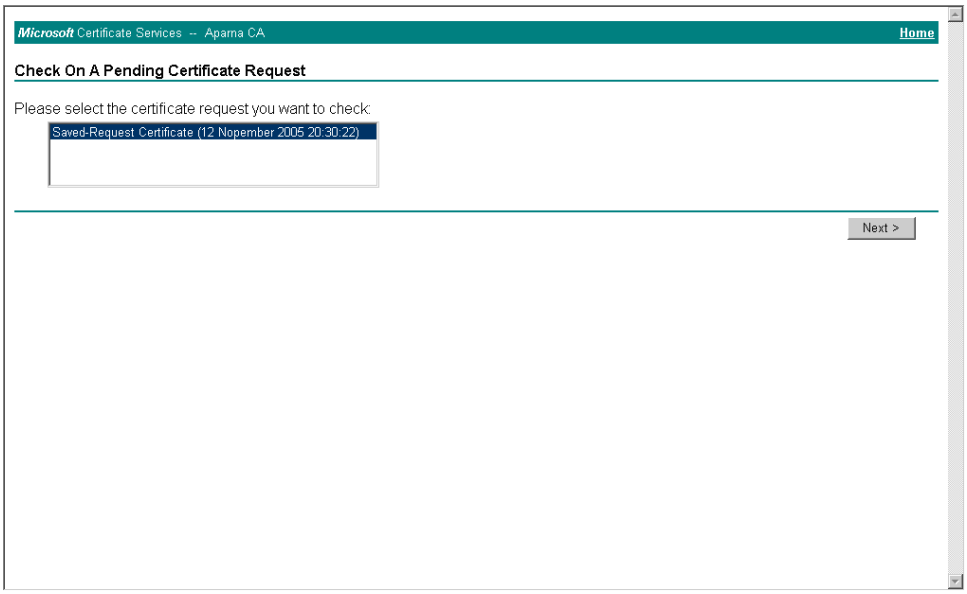

Step 9 Select **Base 64 encoded** and click the **Download CA certificate** link.

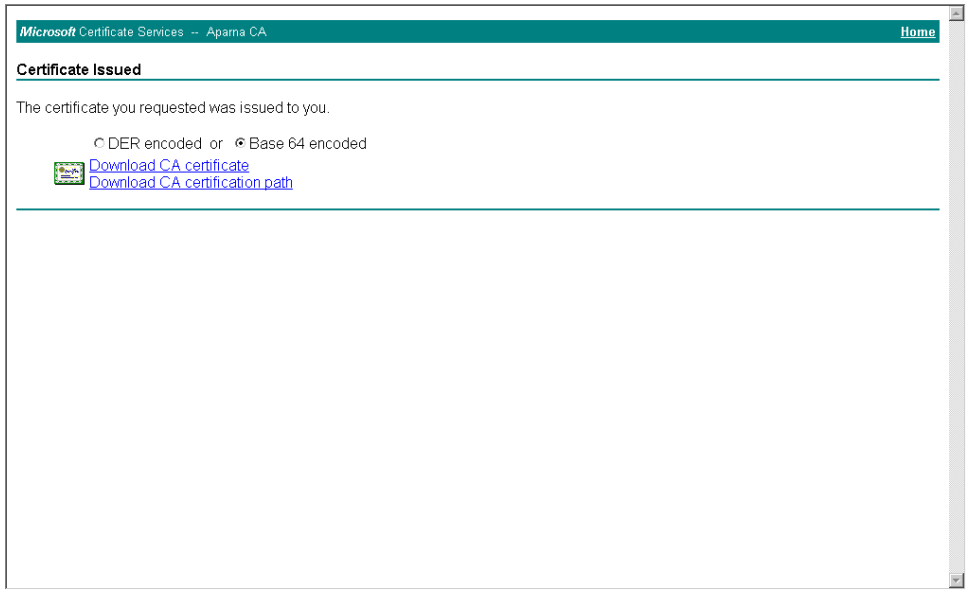

Step 10 Click **Open** on the File Download dialog box.

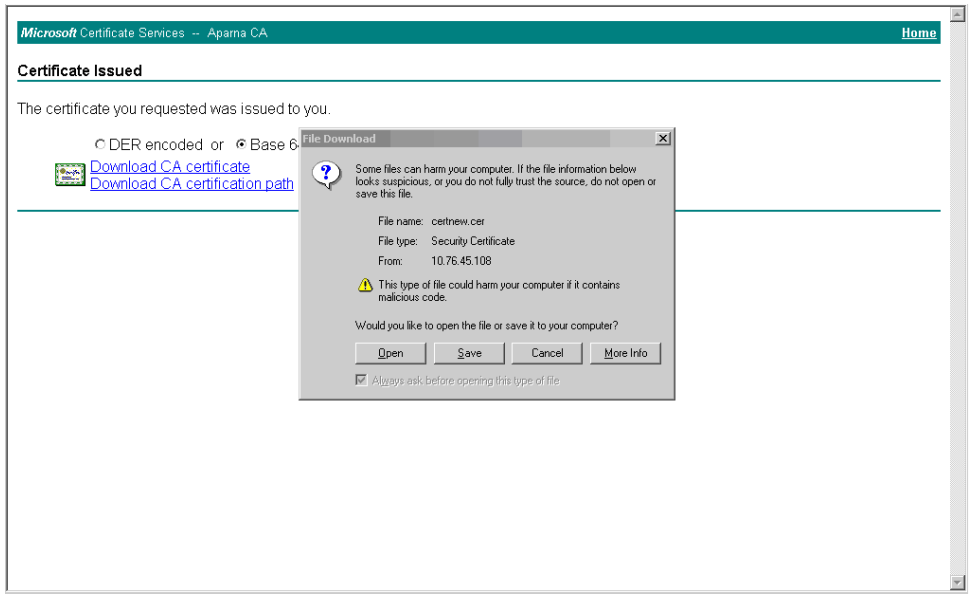

Step 11 Click the **Details** tab on the Certificate dialog and click the **Copy to File** button. Select the **Base-64 encoded X.509 (.CER)** radio button on the Certificate Export Wizard dialog box and click **Next**.

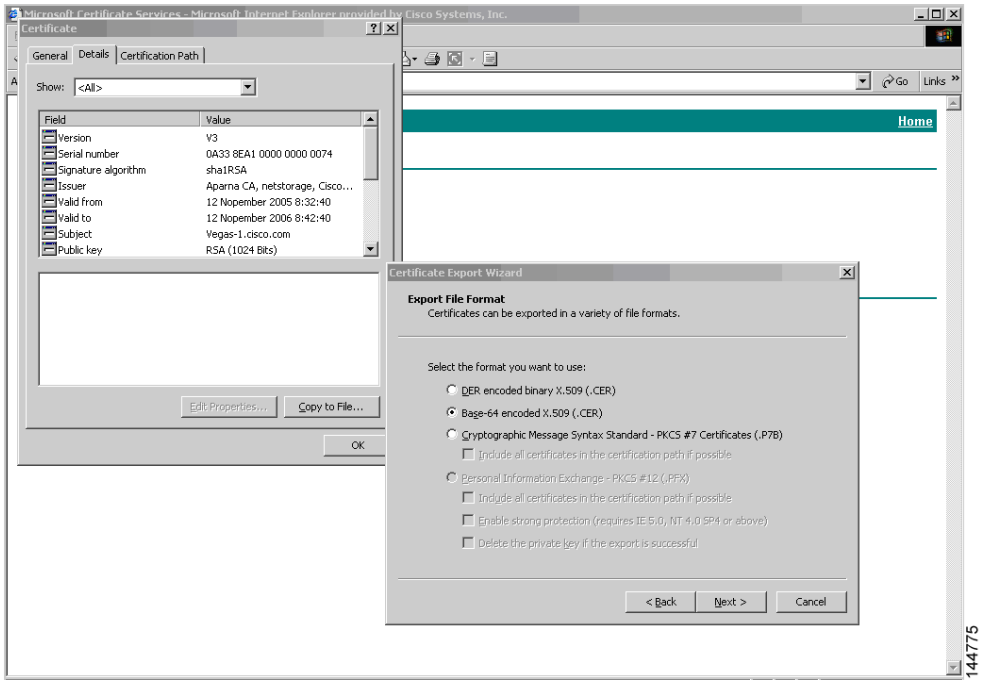

Step 12 Enter the destination file name in the File name: text box on the Certificate Export Wizard dialog box, then click **Next**.

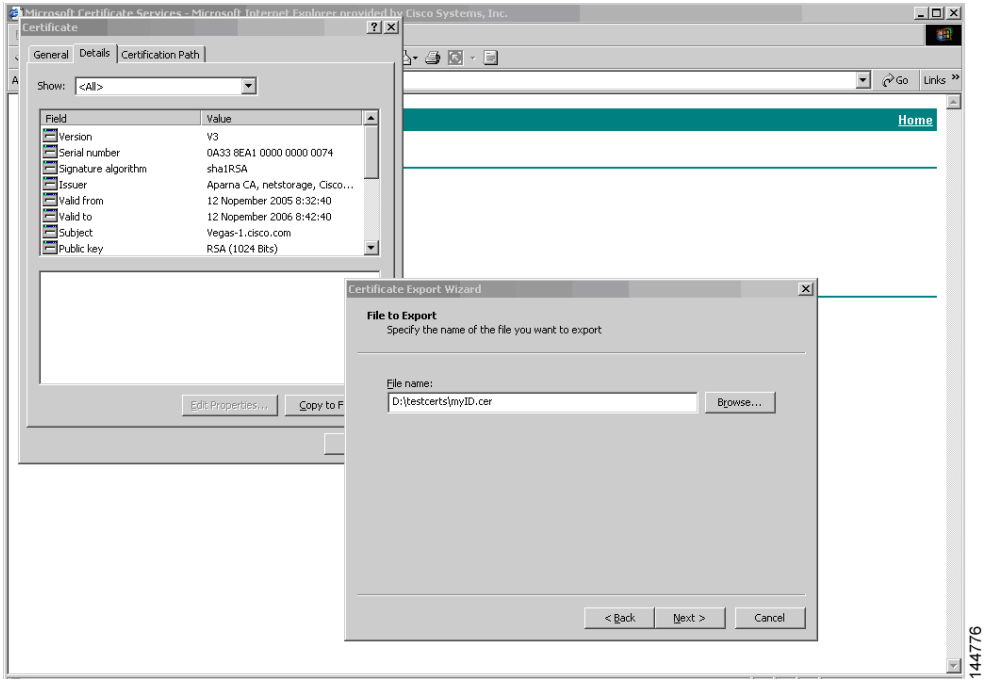

Step 13 Click **Finish**.

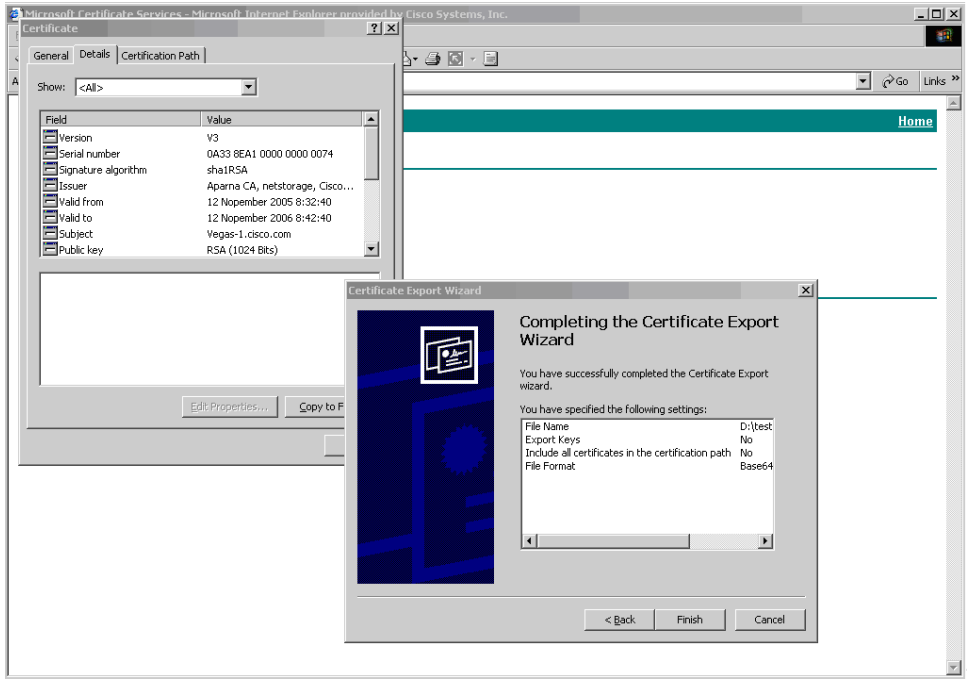

Step 14 Display the identity certificate in base64-encoded format.

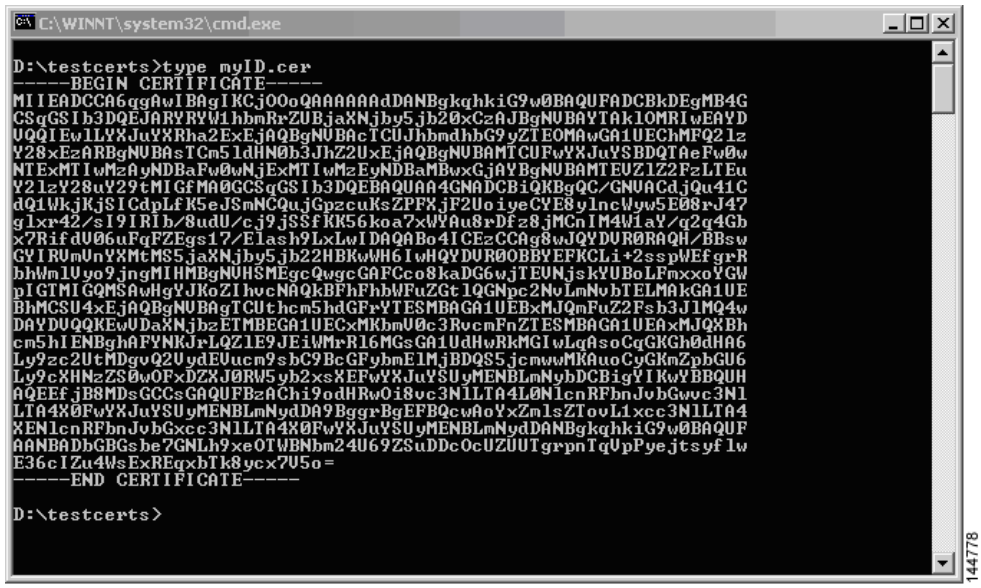

Revoking a Certificate

To revoke a certificate using the Microsoft CA administrator program, follow these steps:

Step 1 Click the **Issued Certificates** folder on the Certification Authority tree. From the list, right-click the certificate you want to revoke.

Step 2 Select **All Tasks > Revoke Certificate**.

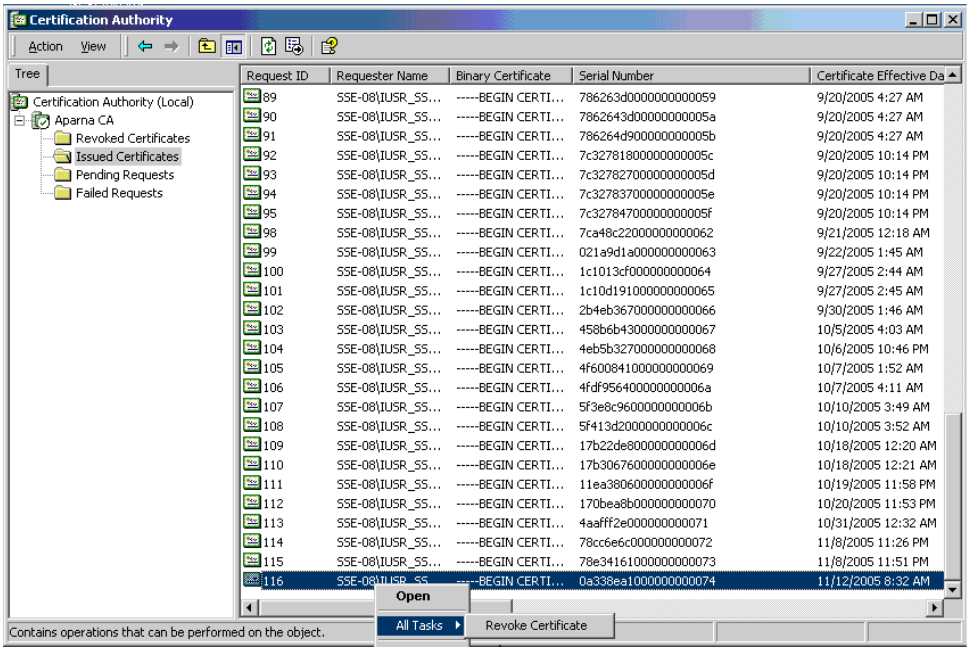

Step 3 Select a reason for the revocation from the Reason code drop-down list, and click **Yes**.

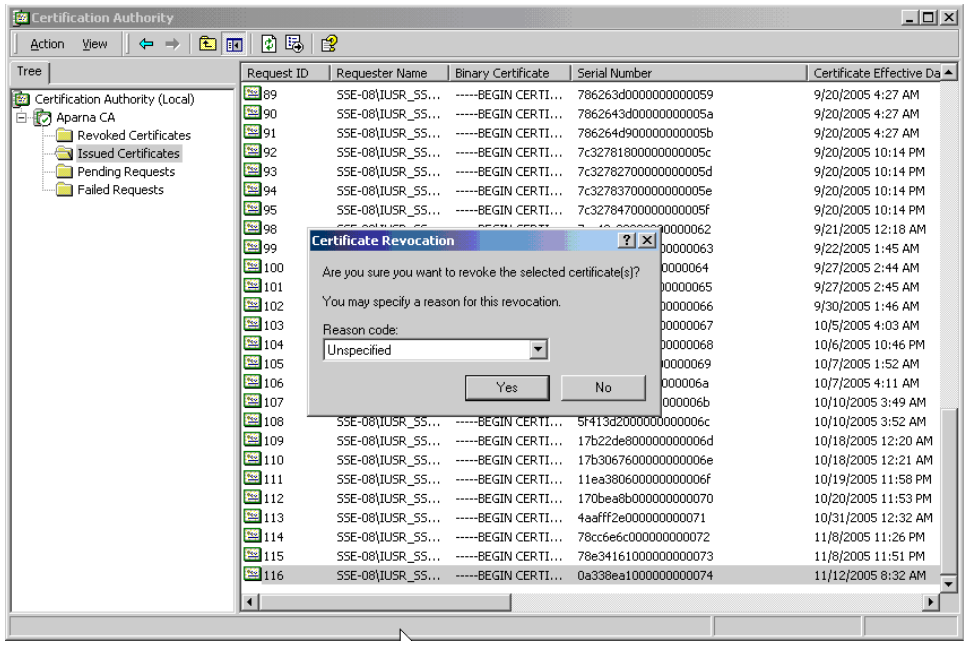

Step 4 Click the **Revoked Certificates** folder to list and verify the certificate revocation.

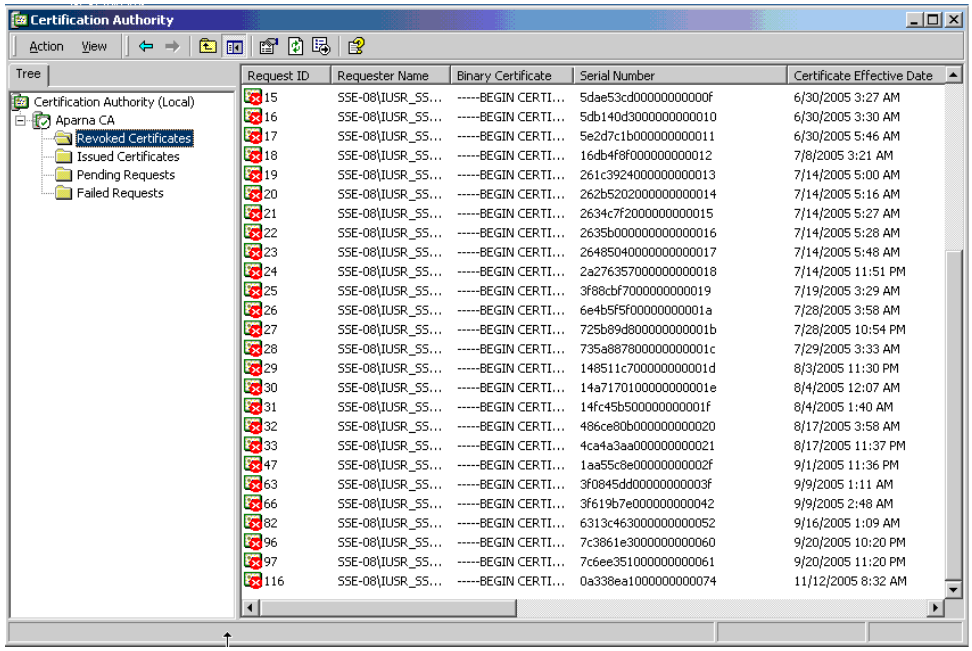

Generating and Publishing the CRL

To generate and publish the CRL using the Microsoft CA administrator program, follow these steps:

Step 1 Select **Action > All Tasks > Publish** on the Certification Authority screen.

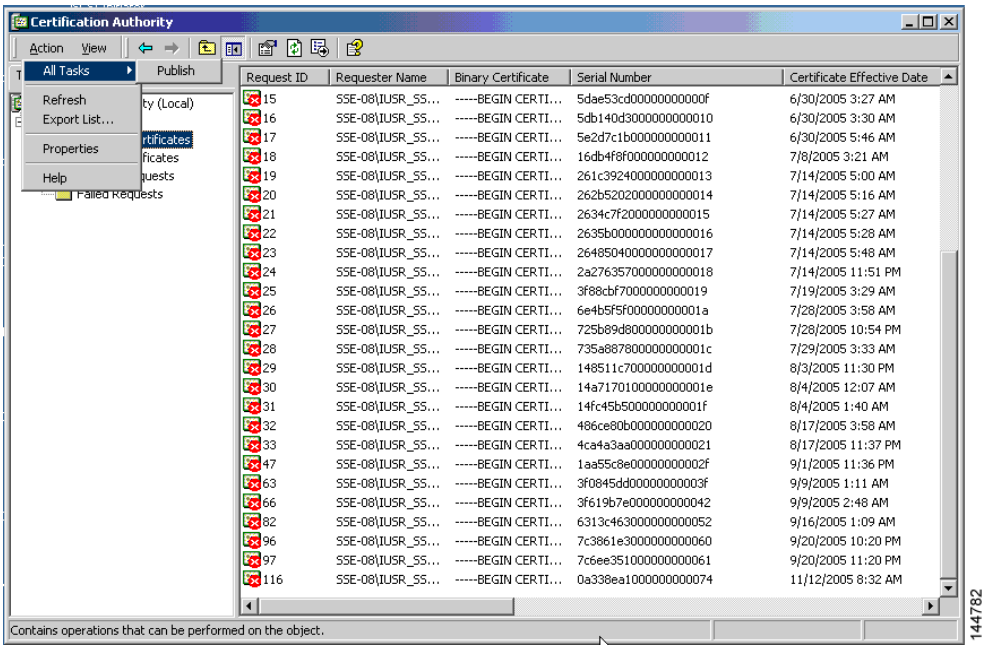

Step 2 Click **Yes** on the Certificate Revocation List dialog box to publish the latest CRL.

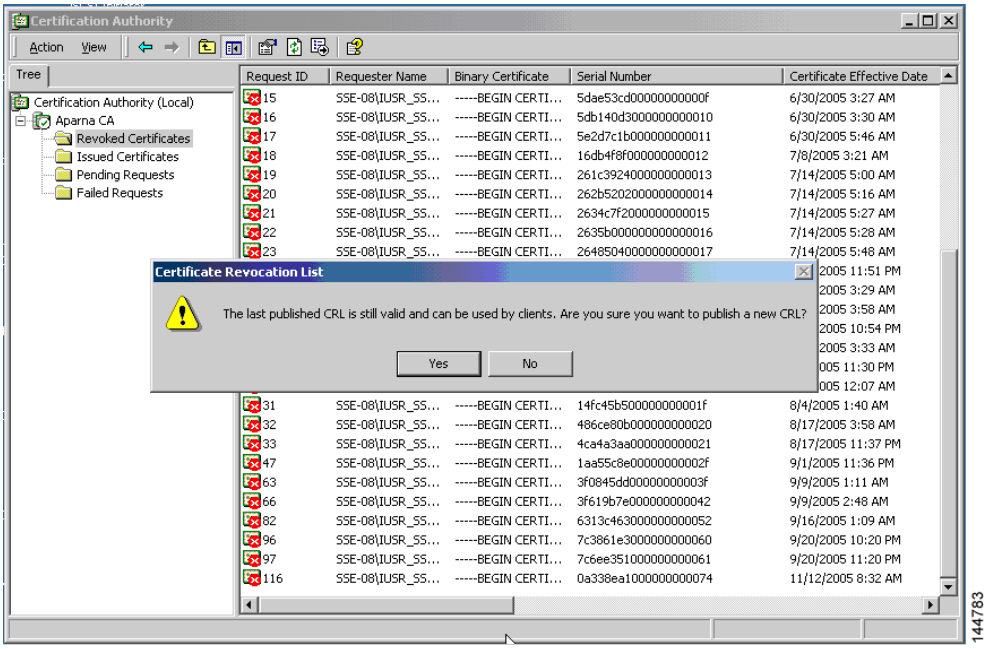

Downloading the CRL

To download the CRL from the Microsoft CA website, follow these steps:

Step 1 Select **Request the CA certificate or certificate revocation list** on the Microsoft Certificate Services web interface and click **Next**.

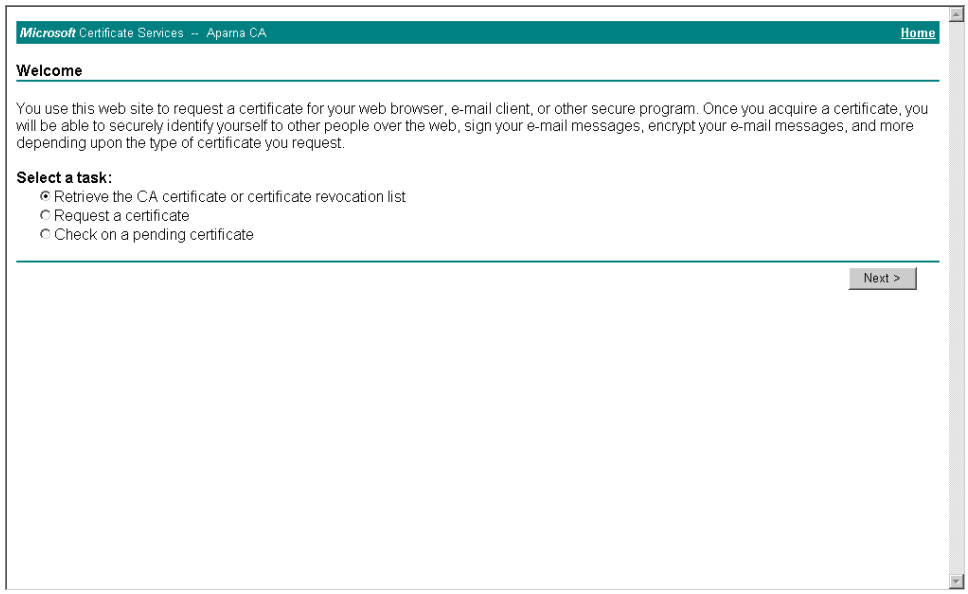

Step 2 Click the **Download latest certificate revocation list** link.

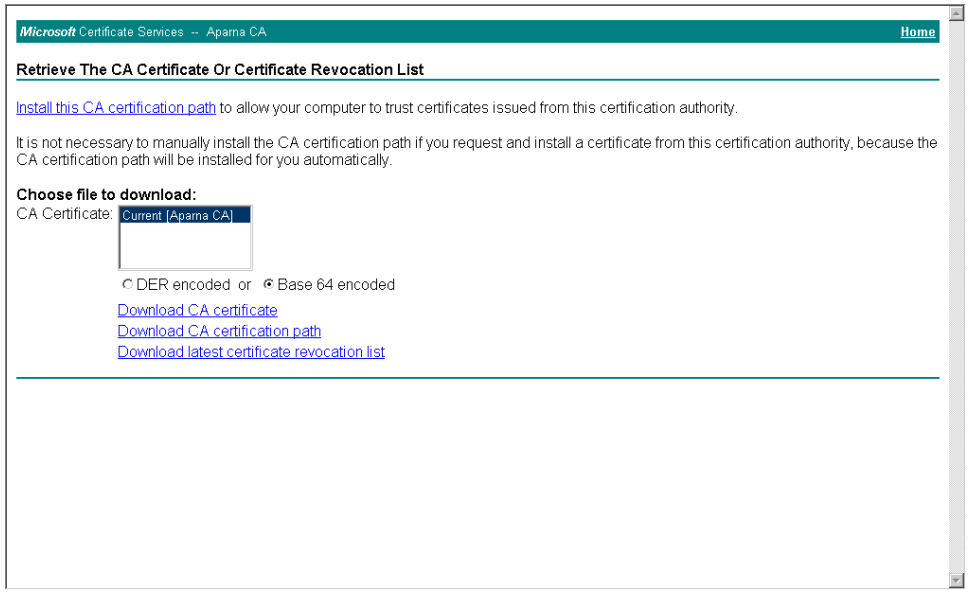

Step 3 Click **Save** in the File Download dialog box.

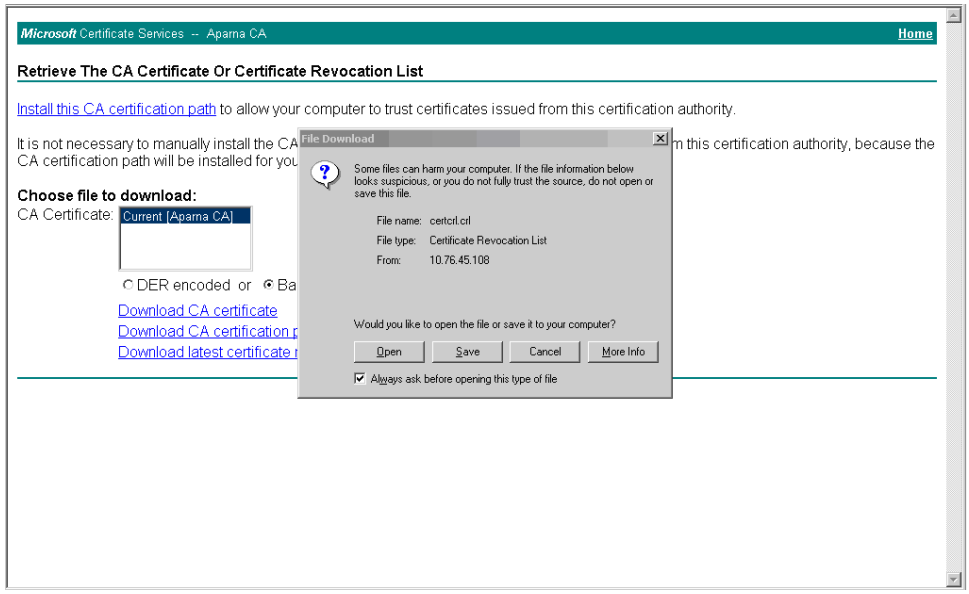

Step 4 Enter the destination file name in the Save As dialog box and click **Save**.

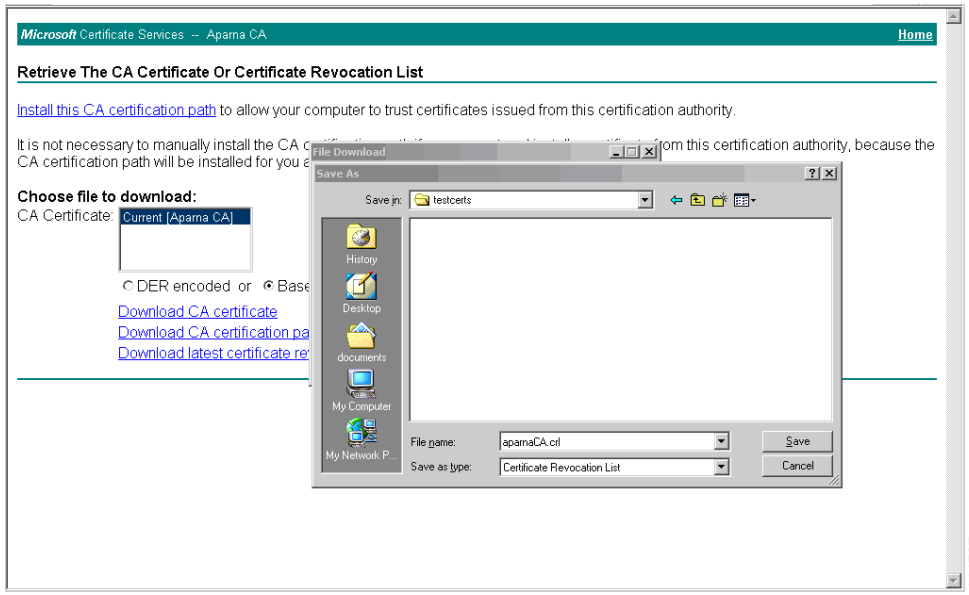

Step 5 Verify the CRL.

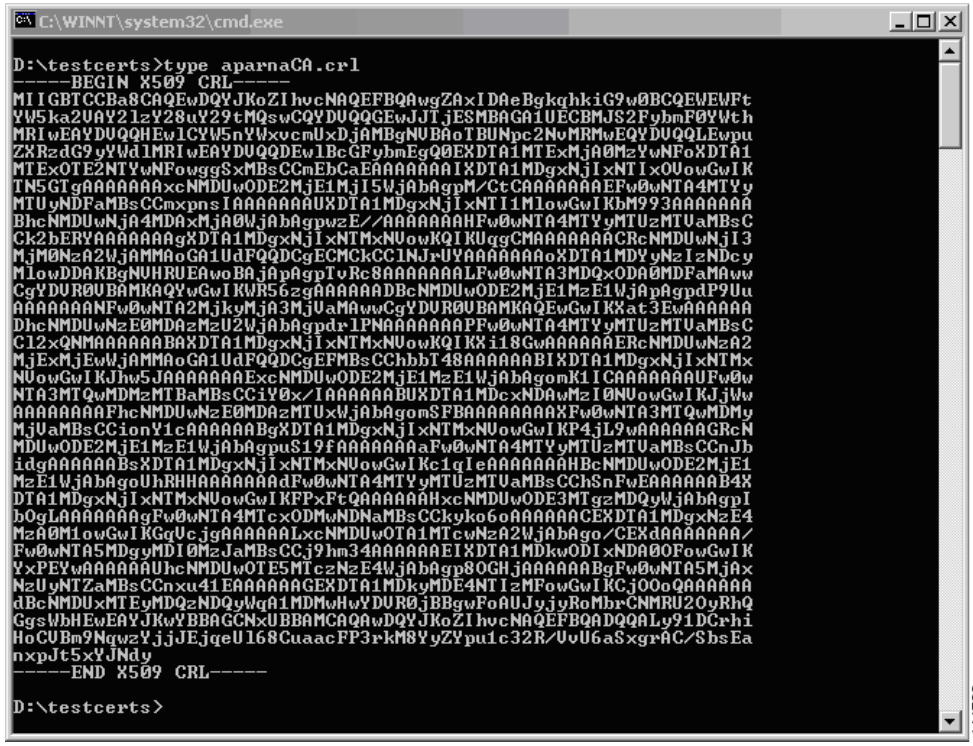

Importing the CRL

To import the CRL to the trust point corresponding to the CA, follow these steps:

Step 1 Click **Switches** > **Security** > **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information pane.

Step 3 Select the **crlimport** option from the Command drop-down menu to import the CRL to the selected trust point.

Step 4 Enter the input file name with the CRL in the bootflash:filename format, in the URL field.

Step 5 Click **Apply Changes** to save the changes.

Note The identity certificate for the switch that was revoked (serial number 0A338EA1000000000074) is listed in the end.

Maximum Limits

Table 38-2 lists the maximum limits for CAs and digital certificate parameters.

Table 38-1 Maximum Limits for CA and Digital Certificate

Feature	Maximum Limit
Trust points declared on a switch	16
RSA key-pairs generated on a switch	16
Identity certificates configured on a switch	16
Certificates in a CA certificate chain	10
Trust points authenticated to a specific CA	10

Default Settings

Table 38-2 lists the default settings for CAs and digital certificate parameters.

Table 38-2 Default CA and Digital Certificate Parameters

Parameters	Default
Trust point	None
RSA key-pair	None
RSA key-pair label	Switch FQDN
RSA key-pair modulus	**512**
RSA key-pair exportable	Yes
Revocation check method of trust point	**crl**

CHAPTER 39

Configuring IPsec Network Security

IP security (IPsec) protocol is a framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. It is developed by the Internet Engineering Task Force (IETF). IPsec provides security services at the IP layer, including protecting one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host. The overall IPsec implementation is per the latest version of RFC 2401. Cisco SAN-OS IPsec implements RFC 2402 through RFC 2410.

IPsec uses the Internet Key Exchange (IKE) protocol to handle protocol and algorithm negotiation and to generate the encryption and authentication keys used by IPsec. While IKE can be used with other protocols, its initial implementation is with the IPsec protocol. IKE provides authentication of the IPsec peers, negotiates IPsec security associations, and establishes IPsec keys. IKE uses RFCs 2408, 2409, 2410, and 2412, and additionally implements the draft-ietf-ipsec-ikev2-16.txt draft.

Note The term IPsec is sometimes used to describe the entire protocol of IPsec data services and IKE security protocols and is other times used to describe only the data services.

This chapter includes the following sections:

About IPsec

IPsec provides security for transmission of sensitive information over unprotected networks such as the Internet. IPsec acts at the network layer, protecting and authenticating IP packets between participating IPsec devices (peers).

IPsec provides the following network security services. In general, the local security policy dictates the use of one or more of these services between two participating IPsec devices:

- Data confidentiality—The IPsec sender can encrypt packets before transmitting them across a network.
- Data integrity—The IPsec receiver can authenticate packets sent by the IPsec sender to ensure that the data has not been altered during transmission.
- Data origin authentication—The IPsec receiver can authenticate the source of the IPsec packets sent. This service is dependent upon the data integrity service.
- Anti-replay protection—The IPsec receiver can detect and reject replayed packets.

Note The term *data authentication* is generally used to mean data integrity and data origin authentication. Within this chapter it also includes anti-replay services, unless otherwise specified.

With IPsec, data can be transmitted across a public network without fear of observation, modification, or spoofing. This enables applications such as Virtual Private Networks (VPNs), including intranets, extranets, and remote user access.

IPsec as implemented in Cisco SAN-OS software supports the Encapsulating Security Payload (ESP) protocol. This protocol encapsulates the data to be protected and provides data privacy services, optional data authentication, and optional anti-replay services.

Note The Encapsulating Security Payload (ESP) protocol is a header inserted into an existing TCP/IP packet, the size of which depends on the actual encryption and authentication algorithms negotiated. To avoid fragmentation, the encrypted packet fits into the interface maximum transmission unit (MTU). The path MTU calculation for TCP takes into account the addition of ESP headers, plus the outer IP header in tunnel mode, for encryption. The MDS switches allow 100 bytes for packet growth for IPsec encryption

Figure 39-1 shows different IPsec scenarios.

Figure 39-1 FCIP and iSCSI Scenarios Using MPS-14/2 Modules

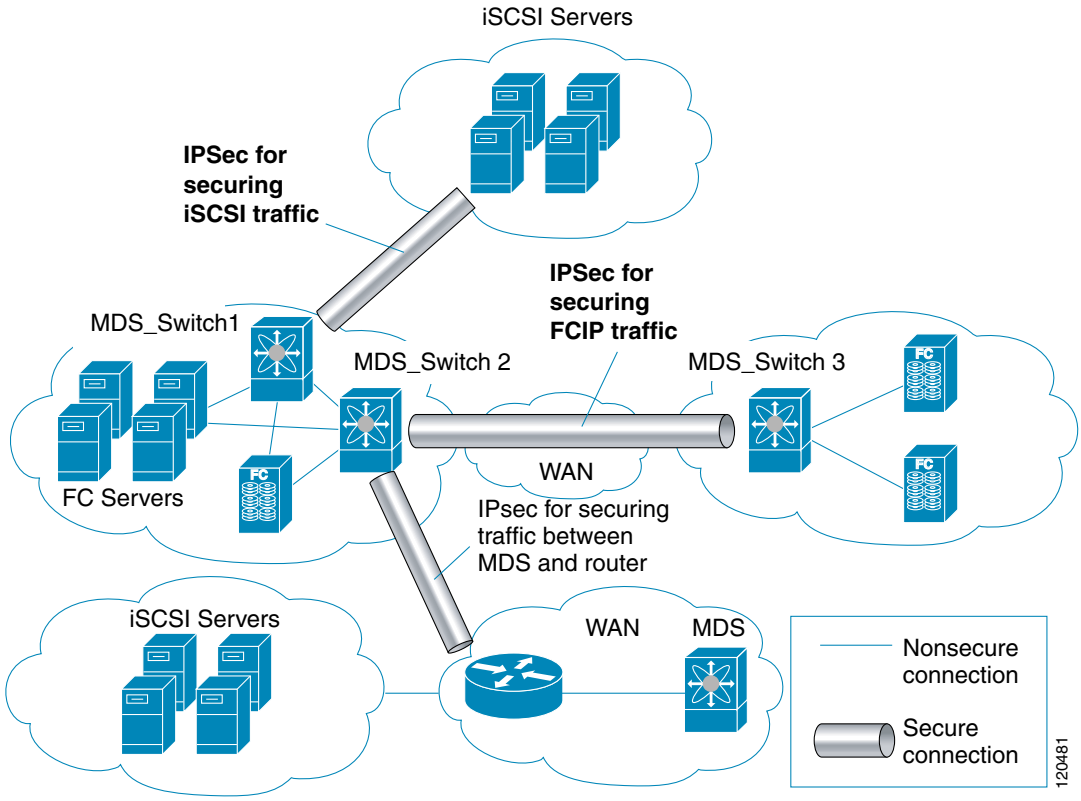

About IKE

IKE automatically negotiates IPsec security associations and generates keys for all switches using the IPsec feature. Specifically, IKE provides these benefits:

- Allows you to refresh IPsec SAs.
- Allows IPsec to provide anti-replay services.
- Supports a manageable, scalable IPsec configuration.
- Allows dynamic authentication of peers.

IPsec Prerequisites

To use the IPsec feature, you need to perform the following tasks:

- Obtain the ENTERPRISE_PKG license (see Chapter 10, "Obtaining and Installing Licenses").
- Configure IKE as described in the "About IKE Initialization" section on page 39-10.

Note The IPsec feature inserts new headers in existing packets (see the "Configuring the MTU Frame Size" section on page 47-6 for more information).

Using IPsec

To use the IPsec feature, follow these steps:

Step 1 Obtain the ENTERPRISE_PKG license to enable IPSEC for iSCSI to enable IPsec for FCIP. See Chapter 10, "Obtaining and Installing Licenses."

Step 2 Configure IKE as described in the "Manually Configuring IPsec and IKE" section on page 39-10.

Note The IPsec feature inserts new headers in existing packets (see the "Configuring the MTU Frame Size" section on page 47-6).

This section contains the following topics:

IPsec Compatibility

IPsec features are compatible with the following Cisco MDS 9000 Family hardware:

- Cisco 14/2-port Multiprotocol Services (MPS-14/2) modules in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors
- Cisco MDS 9216i Switch with the 14/2-port multiprotocol capability in the integrated supervisor module. Refer to the *Cisco MDS 9200 Series Hardware Installation Guide* for more information on the Cisco MDS 9216i Switch.
- The IPsec feature is not supported on the management interface.

IPsec features are compatible with the following fabric setup:

- Two connected Cisco MDS 9200 Switches or Cisco MDS 9500 Directors running Cisco MDS SAN-OS Release 2.0(1b) or later.
- A Cisco MDS 9200 Switches or Cisco MDS 9500 Directors running Cisco MDS SAN-OS Release 2.0(1b) or later connected to any IPsec compliant device.

- The following features are not supported in the Cisco SAN-OS implementation of the IPsec feature:
  - Authentication Header (AH).
  - Transport mode.
  - Security association bundling.
  - Manually configuring security associations.
  - Per host security association option in a crypto map.
  - Security association idle timeout
  - Dynamic crypto maps.

Note Any reference to crypto maps in this document, only refers to static crypto maps.

IPsec and IKE Terminology

The terms used in this chapter are explained in this section.

- Security association (SA)— An agreement between two participating peers on the entries required to encrypt and decrypt IP packets. Two SAs are required for each peer in each direction (inbound and outbound) to establish bidirectional communication between the peers. Sets of bidirectional SA records are stored in the SA database (SAD). IPsec uses IKE to negotiate and bring up SAs. Each SA record includes the following information:
  - Security parameter index (SPI)—A number which, together with a destination IP address and security protocol, uniquely identifies a particular SA. When using IKE to establish the SAs, the SPI for each SA is a pseudo-randomly derived number.
  - Peer—A switch or other device that participates in IPsec. For example, a Cisco MDS switch or other Cisco routers that support IPsec.
  - Transform—A list of operations done to provide data authentication and data confidentiality. For example, one transform is the ESP protocol with the HMAC-MD5 authentication algorithm.
  - Session key—The key used by the transform to provide security services.
  - Lifetime—A lifetime counter (in seconds and bytes) is maintained from the time the SA is created. When the time limit expires the SA is no longer operational and, if required, is automatically renegotiated (rekeyed).
  - Mode of operation—Two modes of operation are generally available for IPsec: tunnel mode and transport mode. The Cisco SAN-OS implementation of IPsec only supports the tunnel mode. The IPsec tunnel mode encrypts and authenticates the IP packet, including its header. The gateways encrypt traffic on behalf of the hosts and subnets.
    The Cisco SAN-OS implementation of IPsec does not support transport mode.

Note The term *tunnel mode* is different from the term *tunnel*, which is used to indicate a secure communication path between two peers, such as two switches connected by an FCIP link.

- Anti-replay—A security service where the receiver can reject old or duplicate packets to protect itself against replay attacks. IPsec provides this optional service by use of a sequence number combined with the use of data authentication.

- Data authentication—Data authentication can refer either to integrity alone or to both integrity and authentication (data origin authentication is dependent on data integrity).
  - Data integrity—Verifies that data has not been altered.
  - Data origin authentication—Verifies that the data was actually sent by the claimed sender.
- Data confidentiality—A security service where the protected data cannot be observed.
- Data flow—A grouping of traffic, identified by a combination of source address and mask or prefix, destination address mask or prefix length, IP next protocol field, and source and destination ports, where the protocol and port fields can have any of these values. Traffic matching a specific combination of these values is logically grouped together into a data flow. A data flow can represent a single TCP connection between two hosts, or it can represent traffic between two subnets. IPsec protection is applied to data flows.
- Perfect forward secrecy (PFS)—A cryptographic characteristic associated with a derived shared secret value. With PFS, if one key is compromised, previous and subsequent keys are not compromised, because subsequent keys are not derived from previous keys.
- Security Policy Database (SPD)—An ordered list of policies applied to traffic. A policy decides if a packet requires IPsec processing, if it should be allowed in clear text, or if it should be dropped.
  - The IPsec SPDs are derived from user configuration of crypto maps.
  - The IKE SPD is configured by the user.

Supported IPsec Transforms and Algorithms

The component technologies implemented for IPsec include the following transforms:

- Advanced Encrypted Standard (AES) is an encryption algorithm. It implements either 128 or 256 bits using Cipher Block Chaining (CBC) or counter mode.
- Data Encryption Standard (DES) is used to encrypt packet data and implements the mandatory 56-bit DES-CBC. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPsec packet.
- Triple DES (3DES) is a stronger form of DES with 168-bit encryption keys that allow sensitive information to be transmitted over untrusted networks.

Note Cisco SAN-OS images with strong encryption are subject to United States government export controls, and have a limited distribution. Images to be installed outside the United States require an export license. Customer orders might be denied or subject to delay due to United States government regulations. Contact your sales representative or distributor for more information, or send e-mail to export@cisco.com.

- Message Digest 5 (MD5) is a hash algorithm with the HMAC variant. HMAC is a keyed hash variant used to authenticate data.
- Secure Hash Algorithm (SHA-1) is a hash algorithm with the Hash Message Authentication Code (HMAC) variant.
- AES-XCBC-MAC is a Message Authentication Code (MAC) using the AES algorithm.

Supported IKE Transforms and Algorithms

The component technologies implemented for IKE include the following transforms:

- Diffie-Hellman (DH) is a public-key cryptography protocol that allows two parties to establish a shared secret over an unsecure communications channel. Diffie-Hellman is used within IKE to establish session keys. Group 1 (768-bit), Group 2 (1024-bit), and Group 5 (1536-bit) are supported.
- Advanced Encrypted Standard (AES) is an encryption algorithm. It implements either 128 bits using Cipher Block Chaining (CBC) or counter mode.
- Data Encryption Standard (DES) is used to encrypt packet data and implements the mandatory 56-bit DES-CBC. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPsec packet.
- Triple DES (3DES) is a stronger form of DES with 168-bit encryption keys that allow sensitive information to be transmitted over untrusted networks.

Note Cisco SAN-OS images with strong encryption are subject to United States government export controls, and have a limited distribution. Images to be installed outside the United States require an export license. Customer orders might be denied or subject to delay due to United States government regulations. Contact your sales representative or distributor for more information, or send e-mail to export@cisco.com.

- Message Digest 5 (MD5) is a hash algorithm with the HMAC variant. HMAC is a keyed hash variant used to authenticate data.
- Secure Hash Algorithm (SHA-1) is a hash algorithm with the Hash Message Authentication Code (HMAC) variant.
- The switch authentication algorithm uses the preshared keys based on the IP address (see "Setting the Default RADIUS Server Timeout Interval and Retransmits" section on page 35-9 for more information on preshared keys).

Configuring IPsec Using FCIP Wizard

Fabric Manager simplifies the configuration of IPsec and IKE by enabling and configuring these features as part of the FCIP configuration using the FCIP Wizard. See the "Using the FCIP Wizard" section on page 43-8.

To enable IPsec using the FCIP Wizard in Fabric Manager, follow these steps:

Step 1 Click the FCIP Wizard icon in the toolbar. (See Figure 39-2.)

Figure 39-2 FCIP Wizard

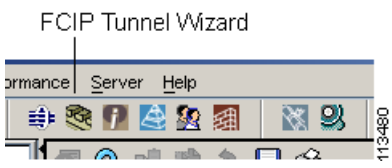

Step 2 Choose the switches that act as endpoints for the FCIP link and click **Next**.

Note These switches must have MPS-14/2 modules installed to configure IPsec on this FCIP link.

Step 3 Choose the Gigabit Ethernet ports on each MPS-14/2 module that will form the FCIP link.

Step 4 Check the **Enforce IPSEC Security** check box and set IKE Auth Key as shown in Figure 39-3.

Figure 39-3 Enabling IPsec on an FCIP Link

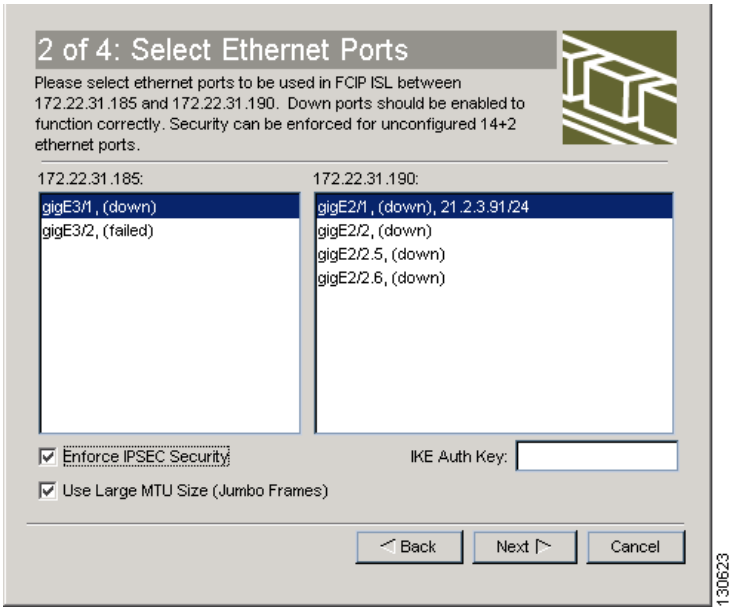

Step 5 Click **Next**. In the Specify Tunnel Properties dialog, you see the TCP connection characteristics.

Step 6 Set the minimum and maximum bandwidth settings and round-trip time for the TCP connections on this FCIP link. Click the **Measure** button to measure the round-trip time between the Gigabit Ethernet endpoints.

Step 7 Check the **Enable Write Acceleration** check box to enable FCIP write acceleration on this FCIP link. See the "FCIP Write Acceleration" section on page 43-23.

Step 8 Check the **Enable Optimum Compression** check box to enable IP compression on this FCIP link. See the "FCIP Compression" section on page 43-30.

Step 9 Click **Next** to configure the FCIP tunnel parameters.

Step 10 Set the Port VSAN for nontrunk/auto and allowed VSAN list for the trunk tunnel. choose a **Trunk Mode** for this FCIP link. See the "Checking Trunk Status" section on page 43-13.

Step 11 Click **Finish** to create this FCIP link or click **Cancel** to exit the FCIP Wizard without creating an FCIP link.

To verify that IPsec and IKE are enabled using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **IPSEC** in the Physical Attributes pane.

You see the IPsec configuration in the Information pane in Figure 39-4.

Figure 39-4 IPSec Configuration

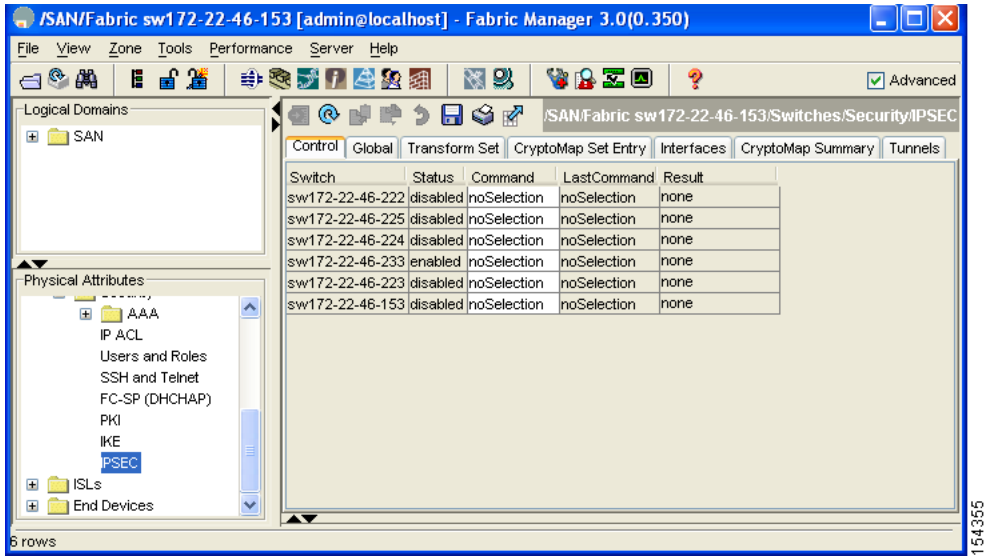

Step 2 The **Control** tab is the default. Verify that the switches you want to modify for IPSec are enabled in the Status column.

Step 3 Expand **Switches > Security** and then select **IKE** in the Physical Attributes pane.

You see the IKE configuration in the Information pane shown in Figure 39-5.

Figure 39-5 IKE Configuration

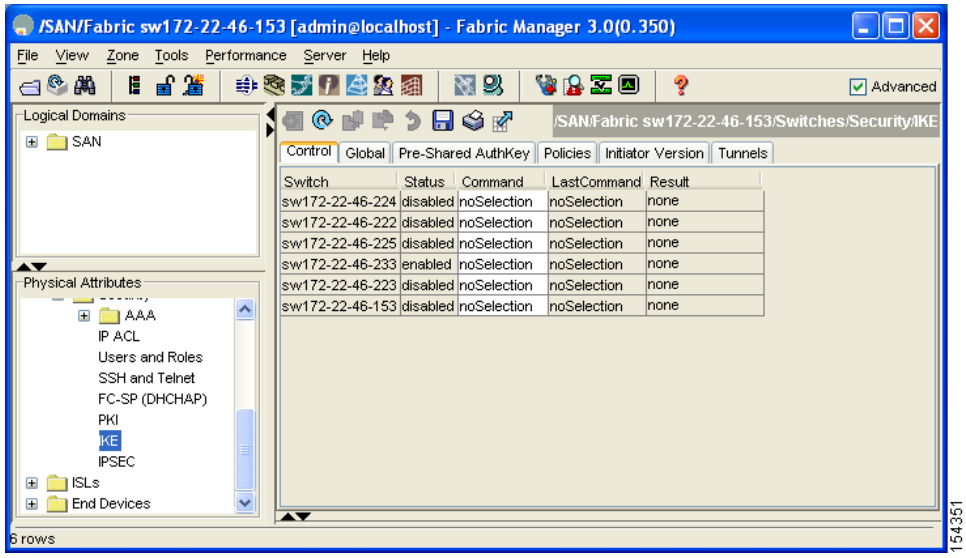

Step 4 The **Control** tab is the default. Verify that the switches you want to modify for IKE are enabled in the Status column.

Manually Configuring IPsec and IKE

This section describes how to manually configure IPsec and IKE if you are not using the FCIP Wizard. See Configuring IPsec Using FCIP Wizard, page 39-7.

IPsec provides secure data flows between participating peers. Multiple IPsec data flows can exist between two peers to secure different data flows, with each tunnel using a separate set of SAs.

After you have completed IKE configuration, configure IPsec.

To configure IPsec in each participating IPsec peer, follow these steps:

Step 1 Identify the peers for the traffic to which secure tunnels should be established.

Step 2 Configure the transform set with the required protocols and algorithms.

Step 3 Create the crypto map and apply access control lists (IPv4 ACLs), transform sets, peers, and lifetime values as applicable.

Step 4 Apply the crypto map to the required interface.

This section contains the following topics:

About IKE Initialization

The IKE feature must first be enabled and configured so the IPsec feature can establish data flow with the required peer. Fabric Manager initializes IKE when you first configure it.

You cannot disable IKE if IPsec is enabled. If you disable the IKE feature, the IKE configuration is cleared from the running configuration.

About the IKE Domain

You must apply the IKE configuration to an IPsec domain to allow traffic to reach the supervisor module in the local switch. Fabric Manager sets the IPsec domain automatically when you configure IKE.

About IKE Tunnels

An IKE tunnel is a secure IKE session between two endpoints. IKE creates this tunnel to protect IKE messages used in IPsec SA negotiations.

Two versions of IKE are used in the Cisco SAN-OS implementation.

- IKE version 1 (IKEv1) is implemented using RFC 2407, 2408, 2409, and 2412.
- IKE version 2 (IKEv2) is a simplified and more efficient version and does not interoperate with IKEv1. IKEv2 is implemented using the draft-ietf-ipsec-ikev2-16.txt draft.

About IKE Policy Negotiation

To protect IKE negotiations, each IKE negotiation begins with a common (shared) IKE policy. An IKE policy defines a combination of security parameters to be used during the IKE negotiation. By default, no IKE policy is configured. You must create IKE policies at each peer. This policy states which security parameters will be used to protect subsequent IKE negotiations and mandates how peers are authenticated. You can create multiple, prioritized policies at each peer to ensure that at least one policy will match a remote peer's policy.

You can configure the policy based on the encryption algorithm (DES, 3DES, or AES), the hash algorithm (SHA or MD5), and the DH group (1, 2, or 5). Each policy can contain a different combination of parameter values. A unique priority number identifies the configured policy. This number ranges from 1 (highest priority) to 255 (lowest priority). You can create multiple policies in a switch. If you need to connect to a remote peer, you must ascertain that at least one policy in the local switch contains the identical parameter values configured in the remote peer. If several policies have identical parameter configurations, the policy with the lowest number is selected.

Table 39-1 provides a list of allowed transform combinations.

Table 39-1 IKE Transform Configuration Parameters

Parameter	Accepted Values	Keyword	Default Value
encryption algorithm	56-bit DES-CBC	**des**	**3des**
	168-bit DES	**3des**	
	128-bit AES	**aes**	
hash algorithm	SHA-1 (HMAC variant)	**sha**	**sha**
	MD5 (HMAC variant)	**md5**	
authentication method	Preshared keys	Not configurable	Preshared keys
DH group identifier	768-bit DH	**1**	**1**
	1024-bit DH	**2**	
	1536-bit DH	**5**	

The following table lists the supported and verified settings for IPsec and IKE encryption authentication algorithms on the Microsoft Windows and Linux platforms:

Platform	IKE	IPsec
Microsoft iSCSI initiator, Microsoft IPsec implementation on Microsoft Windows 2000 platform	3DES, SHA-1 or MD5, DH group 2	3DES, SHA-1
Cisco iSCSI initiator, Free Swan IPsec implementation on Linux platform	3DES, MD5, DH group 1	3DES, MD5

Note When you configure the hash algorithm, the corresponding HMAC version is used as the authentication algorithm.

When the IKE negotiation begins, IKE looks for an IKE policy that is the same on both peers. The peer that initiates the negotiation will send all its policies to the remote peer, and the remote peer will try to find a match. The remote peer looks for a match by comparing its own highest priority policy against the other peer's received policies. The remote peer checks each of its policies in order of its priority (highest priority first) until a match is found.

A match is found when the two peers have the same encryption, hash algorithm, authentication algorithm, and DH group values. If a match is found, IKE completes the security negotiation and the IPsec SAs are created.

If an acceptable match is not found, IKE refuses negotiation and the IPsec data flows will not be established.

Configuring an IKE Policy

To configure the IKE policy negotiation parameters using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **IKE**.

You see the IKE configuration in the Information pane in Figure 39-6.

Figure 39-6 IKE Configuration

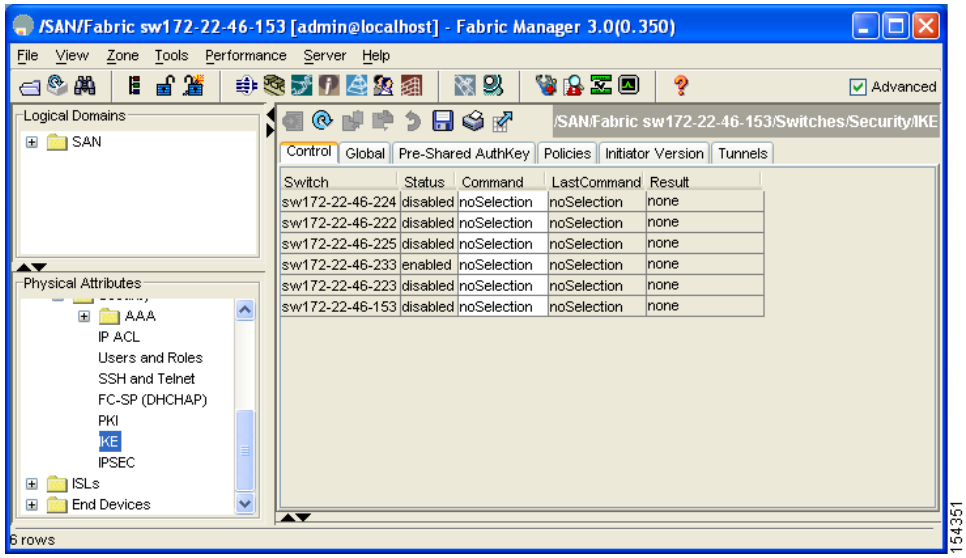

Step 2 Click the **Policies** tab.

You see the existing IKE polices in the Information pane.

Step 3 Click **Create Row** to create an IKE policy.

You see the Create Policy dialog box shown in Figure 39-7.

Figure 39-7 Create IKE

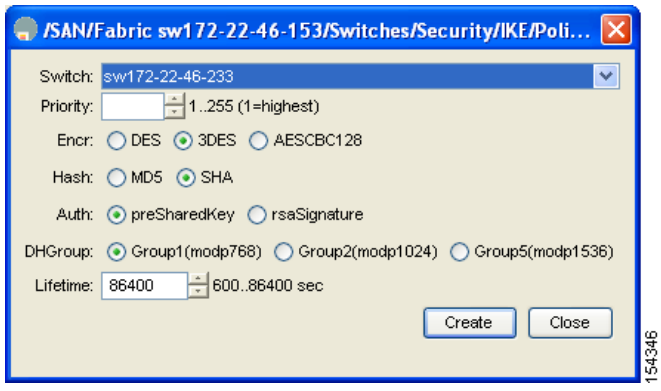

Step 4 Enter the **Priority** for this switch. You can enter a value from one through 255, one being the highest.

Step 5 Select appropriate values for the encryption, hash, authentication, and DHGroup fields.

Step 6 Enter the lifetime for the policy. You can enter a lifetime from 600 to 86400 seconds.

Step 7 Click **Create** to create this policy or click **Close** to discard any unsaved changes (see Figure 39-7).

Note When the authentication method is rsa-sig, make sure the identity hostname is configured for IKE because the IKE certificate has a subject name of the FQDN type.

IPsec Digital Certificate Support

This section describes the advantages of using certificate authorities (CAs) and digital certificates for authentication.

For more information on CAs and digital certificates, see Chapter 38, "Configuring Certificate Authorities and Digital Certificates."

Implementing IPsec Without CAs and Digital Certificates

Without a CA and digital certificates, enabling IPsec services (such as encryption) between two Cisco MDS switches requires that each switch has the key of the other switch (such as an RSA public key or a shared key). You must manually specify either the RSA public keys or preshared keys on each switch in the fabric using IPsec services. Also, each new device added to the fabric will require manual configuration of the other switches in the fabric to support secure communication.

In Figure 39-8, each switch uses the key of the other switch to authenticate the identity of the other switch; this authentication always occurs when IPsec traffic is exchanged between the two switches.

If you have multiple Cisco MDS switches in a mesh topology and wish to exchange IPsec traffic passing among all of those switches, you must first configure shared keys or RSA public keys among all of those switches.

Figure 39-8 Two IPsec Switches Without CAs and Digital Certificates

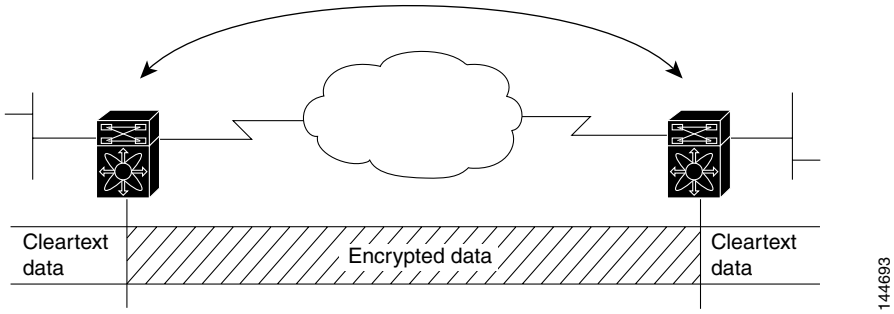

Every time a new switch is added to the IPsec network, you must configure keys between the new switch and each of the existing switches. (In Figure 39-9, four additional two-part key configurations are required to add a single encrypting switch to the network.)

Consequently, the more devices that require IPsec services, the more involved the key administration becomes. This approach does not scale well for larger, more complex encrypting networks.

Figure 39-9 Four IPsec Switches Without a CA and Digital Certificates

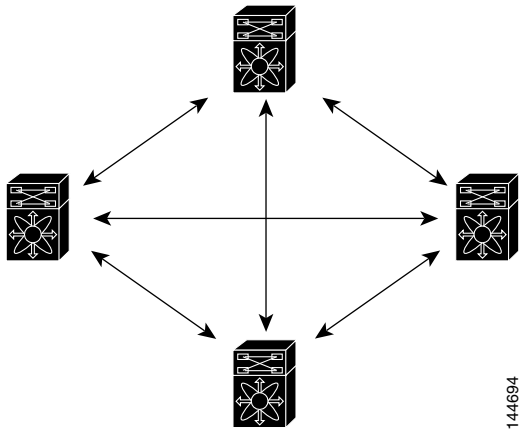

Implementing IPsec with CAs and Digital Certificates

With CA and digital certificates, you do not have to configure keys between all the encrypting switches. Instead, you individually enroll each participating switch with the CA, requesting a certificate for the switch. When this has been accomplished, each participating switch can dynamically authenticate all the other participating switches. When two devices want to communicate, they exchange certificates and digitally sign data to authenticate each other. When a new device is added to the network, you simply enroll that device with a CA, and none of the other devices needs modification. When the new device attempts an IPsec connection, certificates are automatically exchanged and the device can be authenticated.

Figure 39-10 shows the process of dynamically authenticating the devices.

Figure 39-10 Dynamically Authenticating Devices with a CA

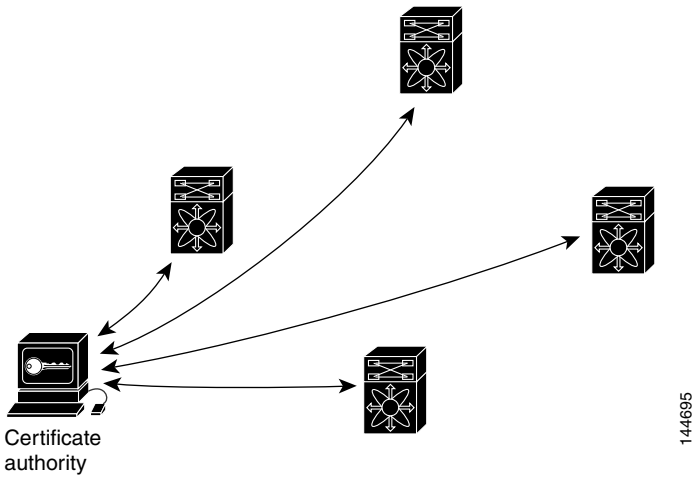

To add a new IPsec switch to the network, you need only configure that new switch to request a certificate from the CA, instead of making multiple key configurations with all the other existing IPsec switches.

How CA Certificates Are Used by IPsec Devices

When two IPsec switches want to exchange IPsec-protected traffic passing between them, they must first authenticate each other—otherwise, IPsec protection cannot occur. The authentication is done with IKE.

IKE can use two methods to authenticate the switches, using preshared keys without a CA and using RSA key-pairs with a CA. Both methods require that keys must be preconfigured between the two switches.

Without a CA, a switch authenticates itself to the remote switch using either RSA-encrypted preshared keys.

With a CA, a switch authenticates itself to the remote switch by sending a certificate to the remote switch and performing some public key cryptography. Each switch must send its own unique certificate that was issued and validated by the CA. This process works because the certificate of each switch encapsulates the public key of the switch, each certificate is authenticated by the CA, and all participating switches recognize the CA as an authenticating authority. This scheme is called IKE with an RSA signature.

Your switch can continue sending its own certificate for multiple IPsec sessions, and to multiple IPsec peers until the certificate expires. When the certificate expires, the switch administrator must obtain a new one from the CA.

CAs can also revoke certificates for devices that will no longer participate in IPsec. Revoked certificates are not recognized as valid by other IPsec devices. Revoked certificates are listed in a certificate revocation list (CRL), which each peer may check before accepting a certificate from another peer.

Certificate support for IKE has the following considerations:

- The switch FQDN (host name and domain name) must be configured before installing certificates for IKE.
- Only those certificates that are configured for IKE or general usage are used by IKE.
- The first IKE or general usage certificate configured on the switch is used as the default certificate by IKE.
- The default certificate is for all IKE peers unless the peer specifies another certificate.
- If the peer asks for a certificate which is signed by a CA that it trusts, then IKE uses that certificate, if it exists on the switch, even if it is not the default certificate.
- If the default certificate is deleted, the next IKE or general usage certificate, if any exists, is used by IKE as the default certificate.
- Certificate chaining is not supported by IKE.
- IKE only sends the identity certificate, not the entire CA chain. For the certificate to be verified on the peer, the same CA chain must also exist there.

Optional IKE Parameter Configuration

You can optionally configure the following parameters for the IKE feature:

- The lifetime association within each policy—The lifetime ranges from 600 to 86,400 seconds. The default is 86,400 seconds (equals one day). The lifetime association within each policy is configured when you are creating an IKE policy. See Configuring an IKE Policy, page 39-13.
- The keepalive time for each peer if you use IKEv2—The keepalive ranges from 120 to 86,400 seconds. The default is 3,600 seconds (equals one hour).
- The initiator version for each peer—IKE v1 or IKE v2 (default). Your choice of initiator version does not affect interoperability when the remote device initiates the negotiation. Configure this option if the peer device supports IKEv1 and you can play the initiator role for IKE with the specified device.
  - If the switches on both sides of an FCIP tunnel are running MDS SAN-OS Release 3.0(1) or later, you must configure initiator version IKEv1 on both sides of an FCIP tunnel to use only IKEv1. If one side of an FCIP tunnel is using IKEv1 and the other side is using IKEv2, the FCIP tunnel uses IKEv2.
  - If the switch on one side of an FCIP tunnel is running MDS SAN-OS Release 3.0(1) or later and the switch on the other side of the FCIP tunnel is running MDS SAN-OS Release 2.x, configuring IKEv1 on either side (or both) results in the FCIP tunnel using IKEv1.

Caution You may need to configure this option even when the switch doesn't behave as an IKE initiator under normal circumstances. Always using this option guarantees a faster recovery of traffic flows in case of failures.

Tip The keepalive time only applies to IKEV2 peers and not to all peers.

Note When IPsec implementations in the host prefer to initiate the IPsec rekey, be sure to configure the IPsec lifetime value in the Cisco MDS switch to be higher than the lifetime value in the host.

This section includes the following topics:

Configuring the Keepalive Time for a Peer

To configure the keepalive time for each peer using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **IKE**.

You see the IKE configuration in the Information pane.

Figure 39-11 IKE Configuration

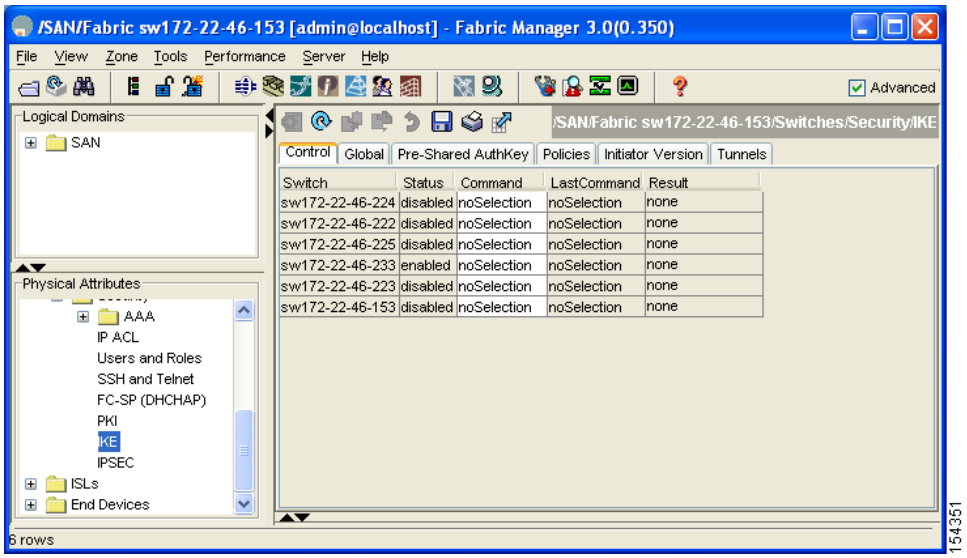

Step 2 Select the **Global** tab.

You see the global statistics of a specific IKE protocol in the Information pane in Figure 39-12.

Figure 39-12 IKE Global Tab Information

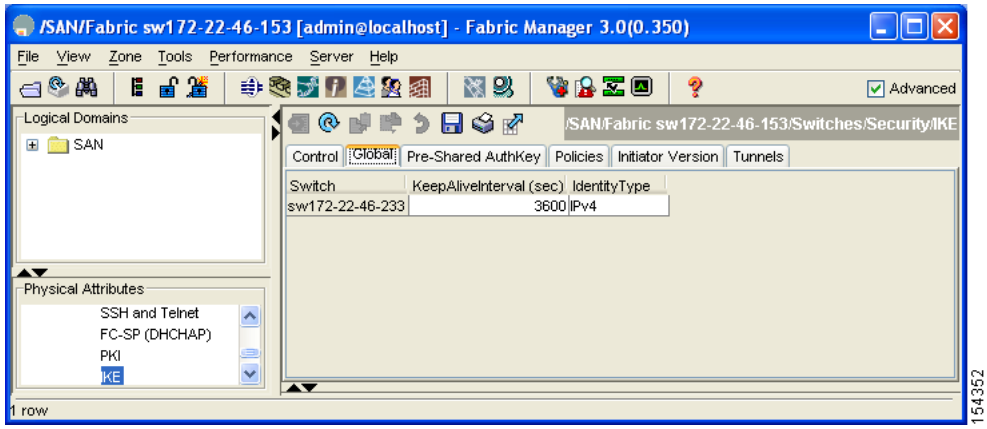

Step 3 Enter a value (in seconds) in the **KeepAliveInterval (sec)**. See Figure 39-12. the The keepalive interval in seconds is used by the IKE entity on the managed device with all the peers for the DOI corresponding to this conceptual row.

Step 4 Click **Apply Changes** to save your changes.

Configuring the Initiator Version

To configure the initiator version, follow these steps:

Step 1 Expand **Switches > Security** and then select **IKE**.

You see the IKE configuration in the Information pane.

Figure 39-13 IKE Configuration

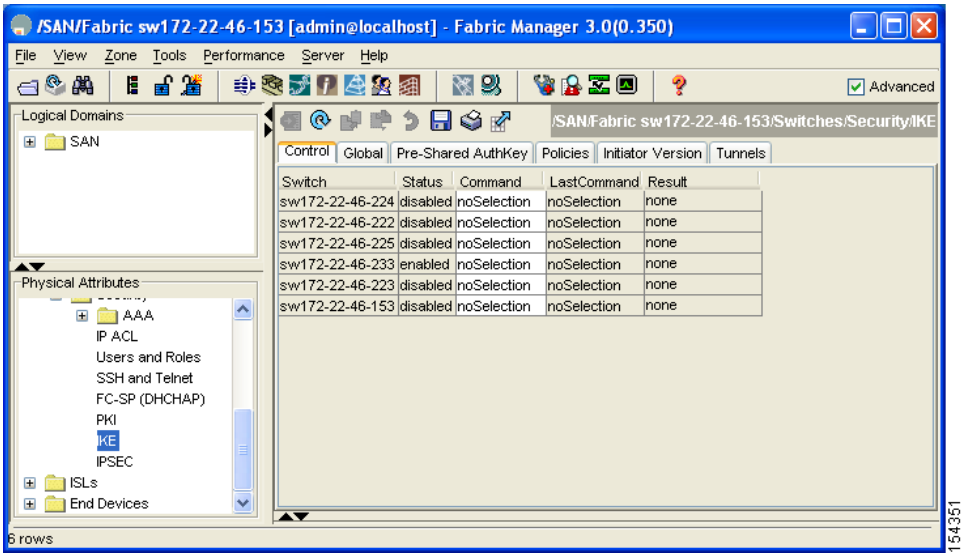

Step 2 Select the **Initiator Version** tab.

You see the existing initiator versions for the peers in the Information pane.

Step 3 Click **Create Row** to create an initiator version.

You see the Create Initiator Version dialog box shown in Figure 39-14.

Figure 39-14 Create Initiator Version Dialog Box

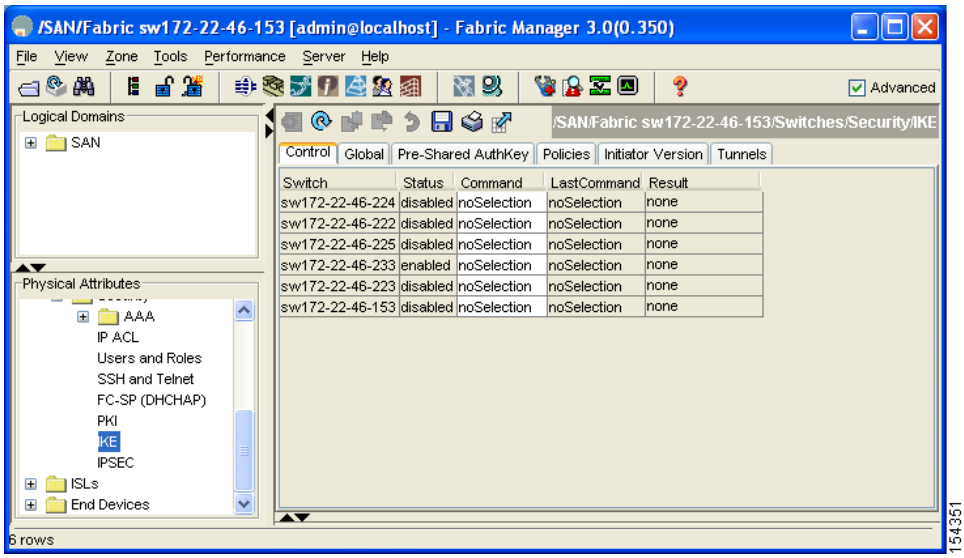

Step 4 Select the Switches for the remote peer for which this IKE protocol initiator is configured.

Step 5 Enter the IP address of the remote peer.

IKEv1 represents the IKE protocol version used when connecting to a remote peer.

Step 6 Click **Create** to create this initiator version or click **Close** to discard any unsaved changes.

Clearing IKE Tunnels or Domains

If an IKE tunnel ID is not specified for the IKE configuration, you can clear all existing IKE domain connections.

To clear all the IKE Tunnels or Domains using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **IKE** in the Physical Attributes pane.

You see the IKE configuration in the Information pane.

Figure 39-15 IKE Configuration

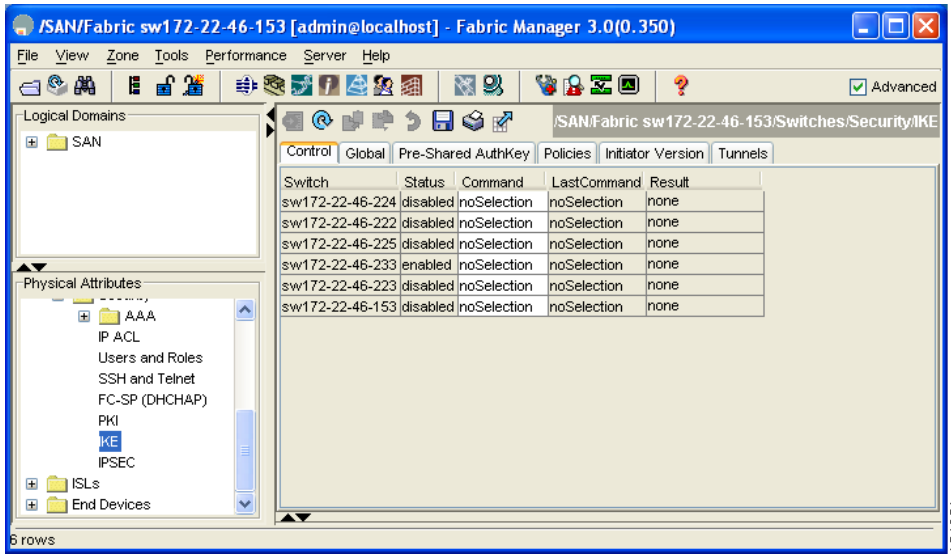

Step 2 Click the **Tunnels** tab in the Information pane.

You see the IKE tunnels.

Step 3 Click the Action column and select **Clear** to clear the tunnel.

Caution When you delete all the SAs within a specific IKEv2 tunnel, then that IKE tunnel is automatically deleted.

Caution When you delete the IKEv2 tunnel, the associated IPsec tunnel under that IKE tunnel is automatically deleted.

Refreshing SAs

To refresh the SAs after changing the IKEv2 configuration using Fabric Manager, follow these steps:

Step 1 Expand **Switches** > **Security** and then select **IKE** in the Physical Attributes pane.

You see the IKE configuration shown in Figure 39-16.

Figure 39-16 IKE Configuration

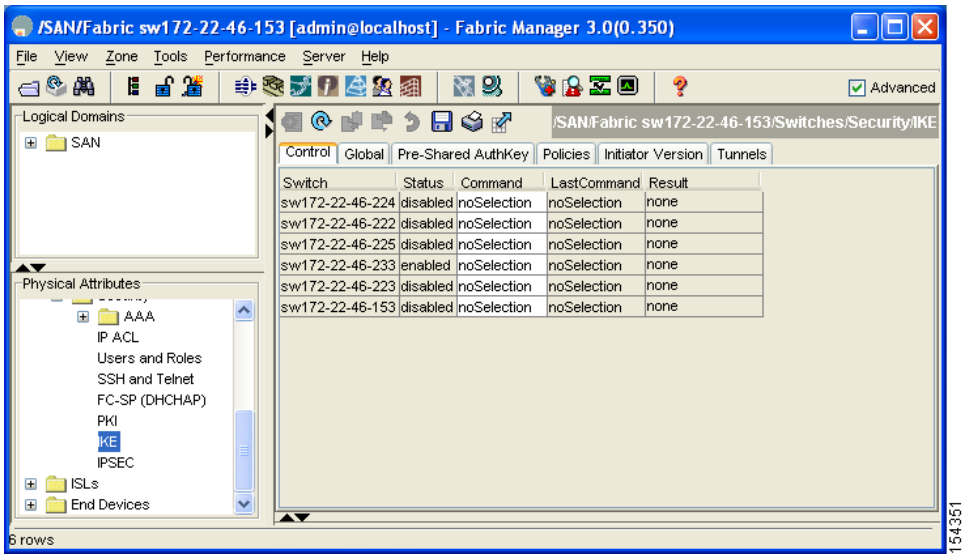

Step 2 Click the **Pre-Shared AuthKey** tab in the Information pane.

Step 3 Click the **Refresh Values**.

Crypto IPv4 ACLs

IP access control lists (IPv4 ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IPv4 ACLs restrict IP-related traffic based on the configured IP filters. See the Chapter 36, "Configuring IPv4 Access Control Lists" for details on creating and defining IPv4 ACLs.

In the context of crypto maps, IPv4 ACLs are different from regular IPv4 ACLs. Regular IPv4 ACLs determine what traffic to forward or block at an interface. For example, IPv4 ACLs can be created to protect all IP traffic between subnet A and subnet Y or Telnet traffic between host A and host B.

This section contains the following topics:

About Crypto IPv4 ACLs

Crypto IPv4 ACLs are used to define which IP traffic requires crypto protection and which traffic does not.

Crypto IPv4 ACLs associated with IPsec crypto map entries have four primary functions:

- Select outbound traffic to be protected by IPsec (permit = protect).
- Indicate the data flow to be protected by the new SAs (specified by a single permit entry) when initiating negotiations for IPsec SAs.
- Process inbound traffic to filter out and discard traffic that should have been protected by IPsec.
- Determine whether or not to accept requests for IPsec SAs on behalf of the requested data flows when processing IKE negotiation from the IPsec peer.

Tip If you want some traffic to receive one type of IPsec protection (for example, encryption only) and other traffic to receive a different type of IPsec protection (for example, both authentication and encryption), create two IPv4 ACLs. Use both IPv4 ACLs in different crypto maps to specify different IPsec policies.

Crypto IPv4 ACL Guidelines

Follow these guidelines when configuring IPv4 ACLs:

- The Cisco SAN-OS software only allows name-based IPv4 ACLs.
- When an IPv4 ACL is applied to a crypto map, the following options apply:
  - Permit—Applies IPv4 ACL.
  - Deny—Allows clear text (default).

  Note IKE traffic (UDP port 500) is implicitly transmitted in clear text.

- The IPsec feature only considers the source and destination IPv4 addresses and subnet masks, protocol, and single port number. There is no support for IPv6 in IPsec.

Note The IPsec feature does not support port number ranges and ignores higher port number field, if specified.

- The permit option causes all IP traffic that matches the specified conditions to be protected by crypto, using the policy described by the corresponding crypto map entry.
- The deny option prevents traffic from being protected by crypto. The first deny statement causes the traffic to be in clear text.
- The crypto IPv4 ACL you define is applied to an interface after you define the corresponding crypto map entry and apply the crypto map set to the interface.
- Different IPv4 ACLs must be used in different entries of the same crypto map set.
- Inbound and outbound traffic is evaluated against the same outbound IPv4 ACL. Therefore, the IPv4 ACL's criteria is applied in the forward direction to traffic exiting your switch, and the reverse direction to traffic entering your switch.
- Each IPv4 ACL filter assigned to the crypto map entry is equivalent to one security policy entry. The IPsec feature supports up to 120 security policy entries for each MPS-14/2 module and Cisco MDS 9216i Switch.
- In Figure 39-17, IPsec protection is applied to traffic between switch interface S0 (IP address 10.0.0.1) and switch interface S1 (IP address 20.0.0.2) as the data exits switch A's S0 interface enroute to switch interface S1. For traffic from 10.0.0.1 to 20.0.0.2, the IPv4 ACL entry on switch A is evaluated as follows:
  - source = IP address 10.0.0.1
  - dest = IP address 20.0.0.2

  For traffic from 20.0.0.2 to 10.0.0.1, that same IPv4 ACL entry on switch A is evaluated as follows:
  - source = IP address 20.0.0.2
  - dest = IP address 10.0.0.1

Figure 39-17 IPsec Processing of Crypto ACLS

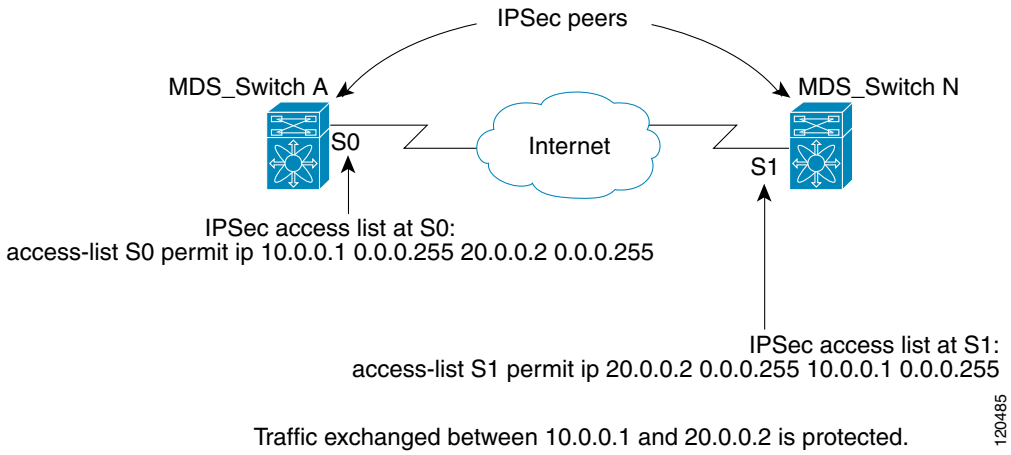

- If you configure multiple statements for a given crypto IPv4 ACL that is used for IPsec, the first permit statement that is matched is used to determine the scope of the IPsec SA. Later, if traffic matches a different permit statement of the crypto IPv4 ACL, a new, separate IPsec SA is negotiated to protect traffic matching the newly matched IPv4 ACL statement.
- Unprotected inbound traffic that matches a permit entry in the crypto IPv4 ACL for a crypto map entry flagged as IPsec is dropped, because this traffic was expected to be protected by IPsec.

Mirror Image Crypto IPv4 ACLs

For every crypto IPv4 ACL specified for a crypto map entry defined at the local peer, define a mirror image crypto IPv4 ACL at the remote peer. This configuration ensures that IPsec traffic applied locally can be processed correctly at the remote peer.

Tip The crypto map entries themselves must also support common transforms and must refer to the other system as a peer.

Figure 39-18 shows some sample scenarios with and without mirror image IPv4 ACLs.

Figure 39-18 IPsec Processing of Mirror Image Configuration

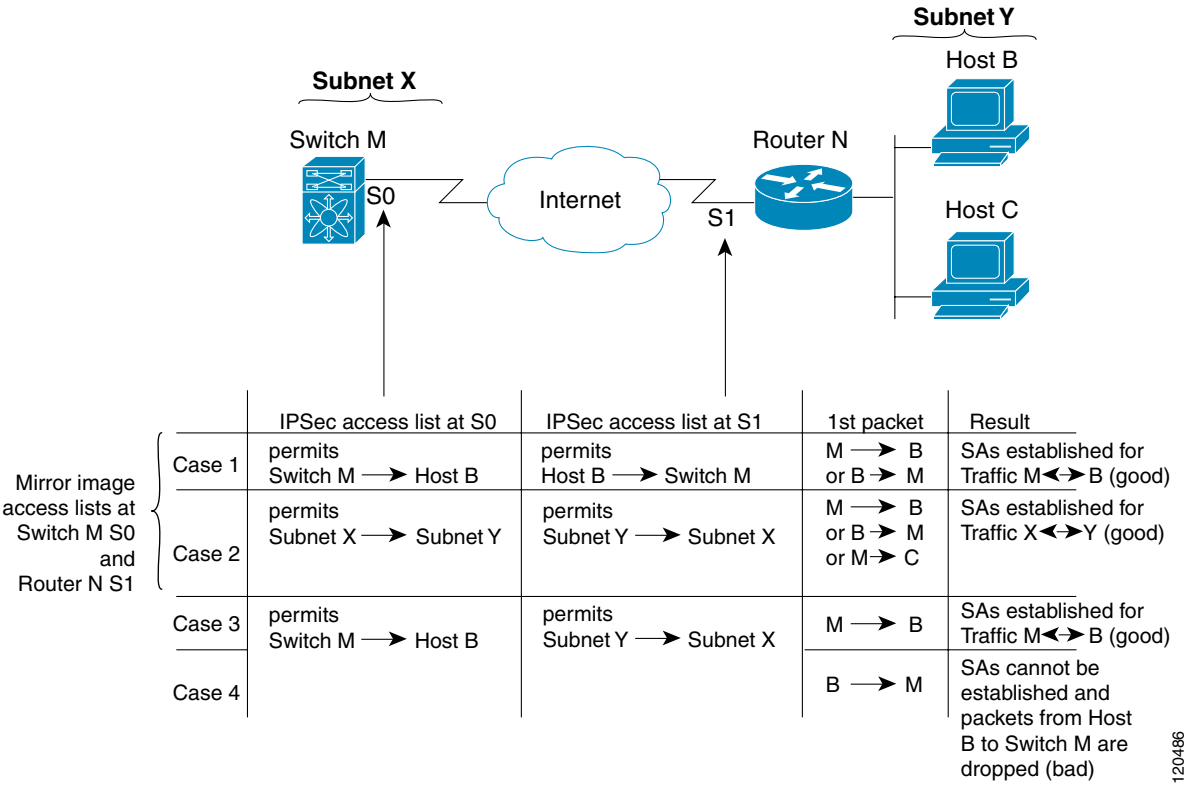

As Figure 39-18 indicates, IPsec SAs can be established as expected whenever the two peers' crypto IPv4 ACLs are mirror images of each other. However, an IPsec SA can be established only some of the time when the IPv4 ACLs are not mirror images of each other. This can happen in the case when an entry in one peer's IPv4 ACL is a subset of an entry in the other peer's IPv4 ACL, as shown in cases 3 and 4.

IPsec SA establishment is critical to IPsec. Without SAs, IPsec does not work, causing any packets matching the crypto IPv4 ACL criteria to be silently dropped instead of being forwarded with IPsec security.

In case 4, an SA cannot be established because SAs are always requested according to the crypto IPv4 ACLs at the initiating packet's end. In case 4, router N requests that all traffic between subnet X and subnet Y be protected, but this is a superset of the specific flows permitted by the crypto IPv4 ACL at switch M so the request is not permitted. Case 3 works because switch M's request is a subset of the specific flows permitted by the crypto IPv4 ACL at router N.

Because of the complexities introduced when crypto IPv4 ACLs are not configured as mirror images at peer IPsec devices, We strongly encourage you to use mirror image crypto IPv4 ACLs.

The any Keyword in Crypto IPv4 ACLs

Tip We recommend that you configure mirror image crypto IPv4 ACLs for use by IPsec and that you avoid using the any option.

The **any** keyword in a permit statement is discouraged when you have multicast traffic flowing through the IPsec interface. This configuration can cause multicast traffic to fail.

The **permit any** statement causes all outbound traffic to be protected (and all protected traffic sent to the peer specified in the corresponding crypto map entry) and requires protection for all inbound traffic. Then, all inbound packets that lack IPsec protection are silently dropped, including packets for routing protocols, NTP, echo, echo response, and so forth.

You need to be sure you define which packets to protect. If you must use **any** in a permit statement, you must preface that statement with a series of deny statements to filter out any traffic (that would otherwise fall within that permit statement) that you do not want to be protected.

Creating Crypto IPv4 ACLs

To create crypto IPv4 ACLs refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

About Transform Sets in IPsec

A transform set represents a certain combination of security protocols and algorithms. During the IPsec security association negotiation, the peers agree to use a particular transform set for protecting a particular data flow.

You can specify multiple transform sets, and then specify one or more of these transform sets in a crypto map entry. The transform set defined in the crypto map entry is used in the IPsec security association negotiation to protect the data flows specified by that crypto map entry's access list.

During IPsec security association negotiations with IKE, the peers search for a transform set that is the same at both peers. When such a transform set is found, it is selected and applied to the protected traffic as part of both peers' IPsec security associations.

Tip If you change a transform set definition, the change is only applied to crypto map entries that reference the transform set. The change is not applied to existing security associations, but used in subsequent negotiations to establish new security associations. If you want the new settings to take effect sooner, you can clear all or part of the security association database.

Note When you enable IPsec, the Cisco SAN-OS software automatically creates a default transform set (ipsec_default_tranform_set) using AES-128 encryption and SHA-1 authentication algorithms.

Table 39-2 provides a list of allowed transform combinations for IPsec.

Table 39-2 IPsec Transform Configuration Parameters

Parameter	Accepted Values	Keyword
encryption algorithm	56-bit DES-CBC	**esp-des**
	168-bit DES	**esp-3des**
	128-bit AES-CBC	**esp-aes 128**
	128-bit AES-CTR[1]	**esp-aes 128 ctr**
	256-bit AES-CBC	**esp-aes 256**
	256-bit AES-CTR[1]	**esp-aes 256 ctr**
hash/authentication algorithm[1] (optional)	SHA-1 (HMAC variant)	**esp-sha1-hmac**
	MD5 (HMAC variant)	**esp-md5-hmac**
	AES-XCBC-MAC	**esp-aes-xcbc-mac**

1. If you configure the AES counter (CTR) mode, you must also configure the authentication algorithm.

The following table lists the supported and verified settings for IPsec and IKE encryption authentication algorithms on the Microsoft Windows and Linux platforms:

Platform	IKE	IPsec
Microsoft iSCSI initiator, Microsoft IPsec implementation on Microsoft Windows 2000 platform	3DES, SHA-1 or MD5, DH group 2	3DES, SHA-1
Cisco iSCSI initiator, Free Swan IPsec implementation on Linux platform	3DES, MD5, DH group 1	3DES, MD5

Configuring Transform Sets

To configure transform sets using Fabric Manager, follow these steps:

Step 1 Expand **Switches** > **Security** and then select **IPSec** in the Physical Attributes pane.

You see the IPSec configuration in Figure 39-19.

Figure 39-19 IPsec Configuration

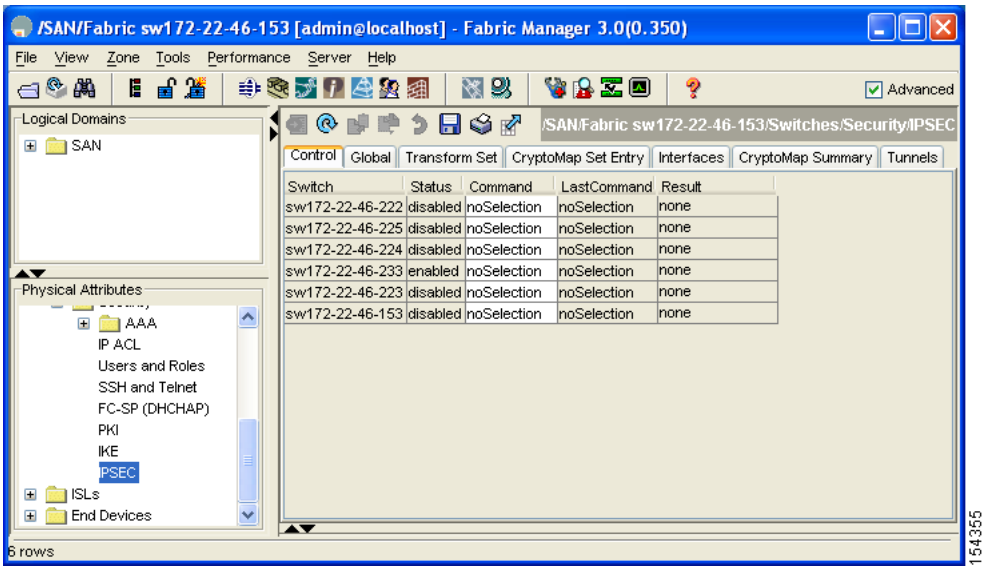

Step 2 Click the **Transform Set** tab in the Information pane.

Step 3 Click **Create Row**.

You see the Create IPSEC dialog box. See Figure 39-20.

Figure 39-20 Create IPSEC

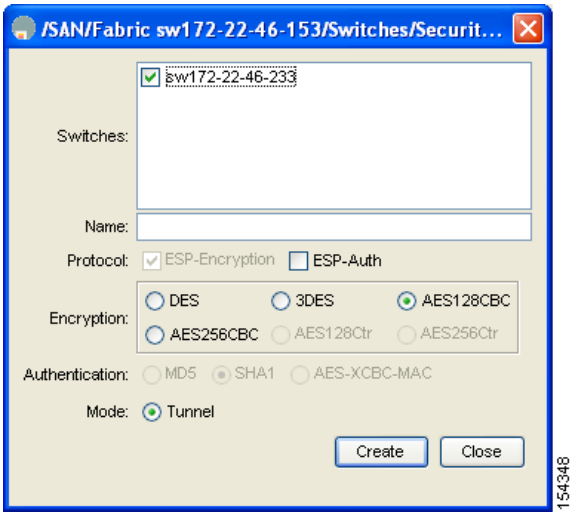

Step 4 Select the switches that you want to create a transform set for in the Create Transform Set dialog box.

Step 5 Assign a name and protocol for the transform set.

Step 6 Select the encryption and authentication algorithm. See Table 39-2 to verify the allowed transform combinations.

Step 7 Click **Create** to create the transform set or you click **Close**.

About Crypto Map Entries

Once you have created the crypto IPv4 ACLs and transform sets, you can create crypto map entries that combine the various parts of the IPsec SA, including the following:

- The traffic to be protected by IPsec (per the crypto IPv4 ACL). A crypto map set can contain multiple entries, each with a different IPv4 ACL.
- The granularity of the flow to be protected by a set of SAs.
- The IPsec-protected traffic destination (who the remote IPsec peer is).
- The local address to be used for the IPsec traffic (applying to an interface).
- The IPsec security to be applied to this traffic (selecting from a list of one or more transform sets).
- Other parameters to define an IPsec SA.

Crypto map entries with the same crypto map name (but different map sequence numbers) are grouped into a crypto map set.

When you apply a crypto map set to an interface, the following events occur:

- A security policy database (SPD) is created for that interface.
- All IP traffic passing through the interface is evaluated against the SPD.

If a crypto map entry sees outbound IP traffic that requires protection, an SA is negotiated with the remote peer according to the parameters included in the crypto map entry.

The policy derived from the crypto map entries is used during the negotiation of SAs. If the local switch initiates the negotiation, it will use the policy specified in the crypto map entries to create the offer to be sent to the specified IPsec peer. If the IPsec peer initiates the negotiation, the local switch checks the policy from the crypto map entries and decide whether to accept or reject the peer's request (offer).

For IPsec to succeed between two IPsec peers, both peers' crypto map entries must contain compatible configuration statements.

SA Establishment Between Peers

When two peers try to establish an SA, they must each have at least one crypto map entry that is compatible with one of the other peer's crypto map entries.

For two crypto map entries to be compatible, they must at least meet the following criteria:

- The crypto map entries must contain compatible crypto IPv4 ACLs (for example, mirror image IPv4 ACLs). If the responding peer entry is in the local crypto, the IPv4 ACL must be permitted by the peer's crypto IPv4 ACL.
- The crypto map entries must each identify the other peer or must have auto peer configured.

- If you create more than one crypto map entry for a given interface, use the `seq-num` of each map entry to rank the map entries: the lower the `seq-num`, the higher the priority. At the interface that has the crypto map set, traffic is evaluated against higher priority map entries first.
- The crypto map entries must have at least one transform set in common where IKE negotiations are carried out and SAs are established. During the IPsec SA negotiation, the peers agree to use a particular transform set when protecting a particular data flow.

When a packet matches a permit entry in a particular IPv4 ACL, the corresponding crypto map entry is tagged, and connections are established.

Crypto Map Configuration Guidelines

When configuring crypto map entries, follow these guidelines:

- The sequence number for each crypto map decides the order in which the policies are applied. A lower sequence number is assigned a higher priority.
- Only one IPv4 ACL is allowed for each crypto map entry (the IPv4 ACL itself can have multiple permit or deny entries).
- When the tunnel endpoint is the same as the destination address, you can use the auto-peer option to dynamically configure the peer.
- For IPsec to interoperate effectively with Microsoft iSCSI initiators, specify the TCP protocol and the local iSCSI TCP port number (default 3260) in the IPv4-ACL. This configuration ensures the speedy recovery of encrypted iSCSI sessions following disruptions such as Gigabit Ethernet interfaces shutdowns, VRRP switchovers, and port failures.

Creating Crypto Map Entries

To create mandatory crypto map entries using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **IPSEC** in the Physical Attributes pane.

You see the IPSec configuration in the Information pane in Figure 39-21.

Figure 39-21 IPsec Configuration

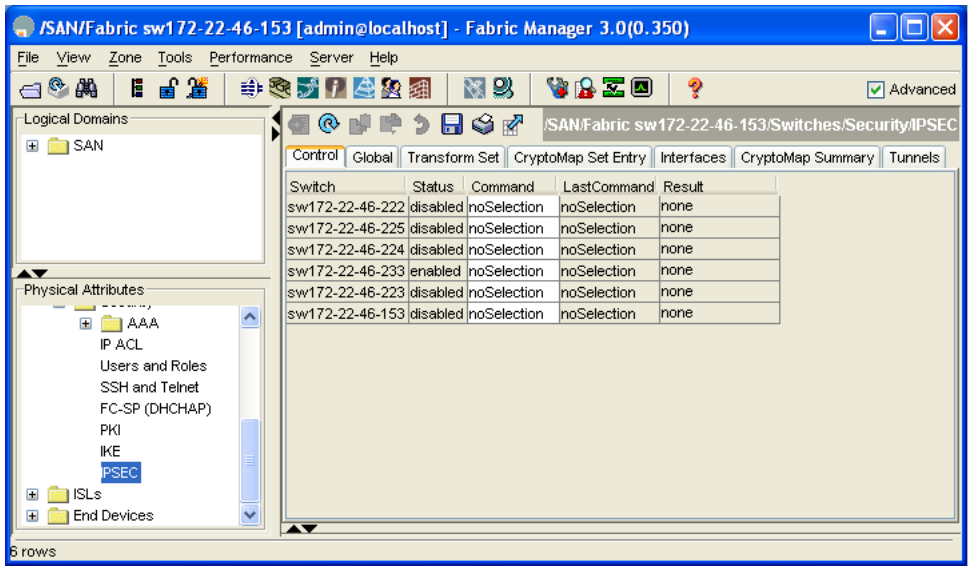

Step 2 Choose the **CryptoMap Set Entry** tab.

You see the existing crypto maps configured in Figure 39-22.

Figure 39-22 Existing Crypto Maps

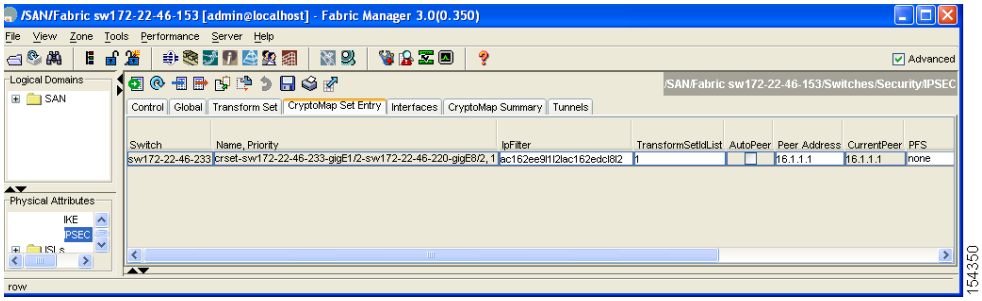

Step 3 Optionally, click **Create Row** to create a crypto map entry.

You see the Create Crypto Map dialog box. See Figure 39-23.

Figure 39-23 Create Crypto Map Dialog Box

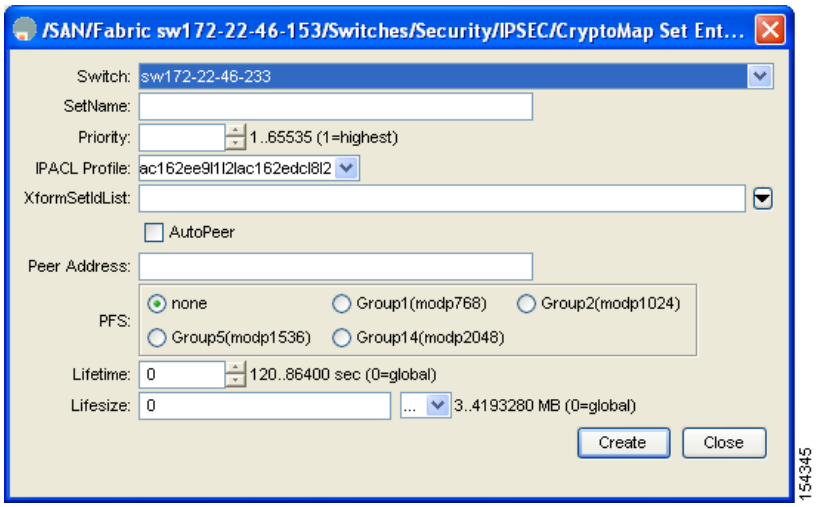

Step 4 Select the switch that you want to configure or modify. If you are creating a crypto map, set the setName and priority for this crypto map.

Step 5 Select the IPv4 ACL Profile and TransformSetIdList from the drop-down list for this crypto map.

Step 6 Optionally, check the **AutoPeer** check box or set the peer address if you are creating a crypto map. See the "About the AutoPeer Option" section on page 39-34.

Step 7 Choose the appropriate PFS selection. See the "About Perfect Forward Secrecy" section on page 39-37.

Step 8 Supply the Lifetime and LifeSize. See the "About SA Lifetime Negotiation" section on page 39-32.

Step 9 Click **Create** if you are creating a crypto map, or click **Apply Changes** if you are modifying an existing crypto map.

About SA Lifetime Negotiation

You can override the global lifetime values (size and time) by configuring an SA-specific lifetime value.

To specify SA lifetime negotiation values, you can optionally configure the lifetime value for a specified crypto map. If you do, this value overrides the globally set values. If you do not specify the crypto map specific lifetime, the global value (or global default) is used.

See the "Global Lifetime Values" section on page 39-39 for more information on global lifetime values.

Setting the SA Lifetime

To set the SA lifetime for a specified crypto map entry using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **IPSEC** in the Physical Attributes pane.

You see the IP SEC configuration in the Information pane.

Figure 39-24 IPsec Configuration

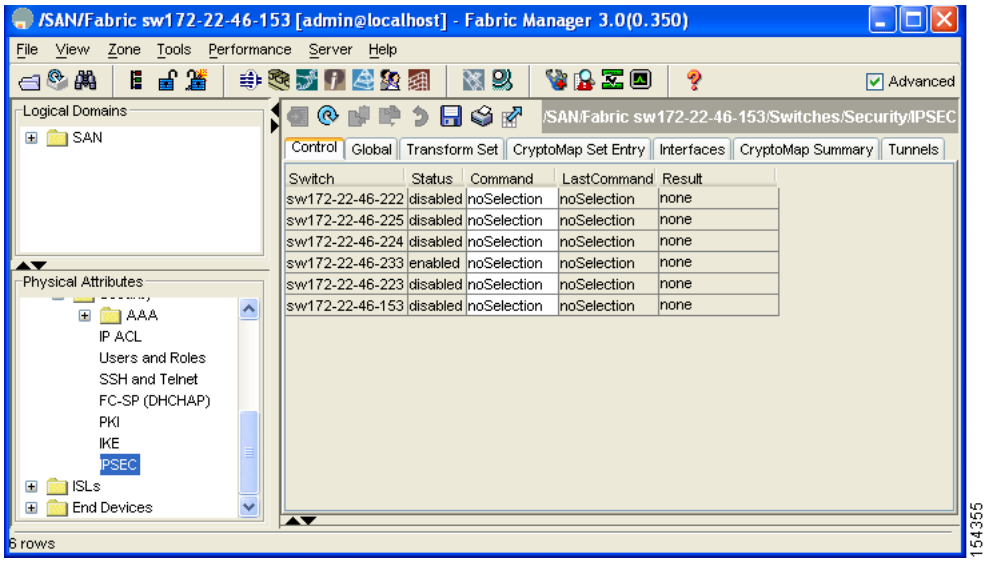

Step 2 Choose the **CryptoMap Set Entry** tab.

You see the existing crypto maps configured in Figure 39-25.

Figure 39-25 Existing Crypto Maps - Leftmost Columns

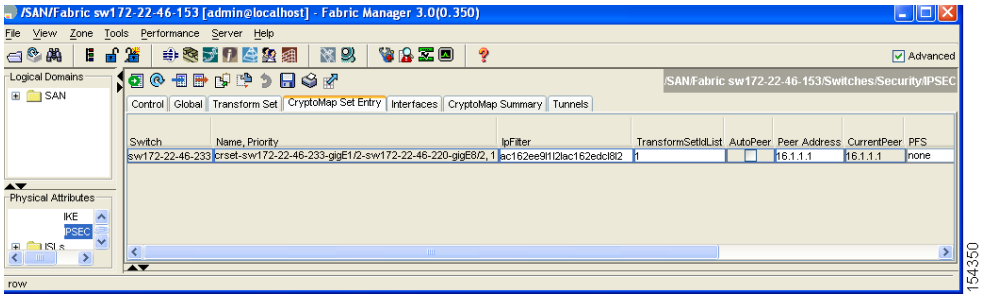

Step 3 Scroll to the right half of the dialog box.

You see more columns shown in Figure 39-26.

Figure 39-26 Existing Crypto Maps - Rightmost Columns

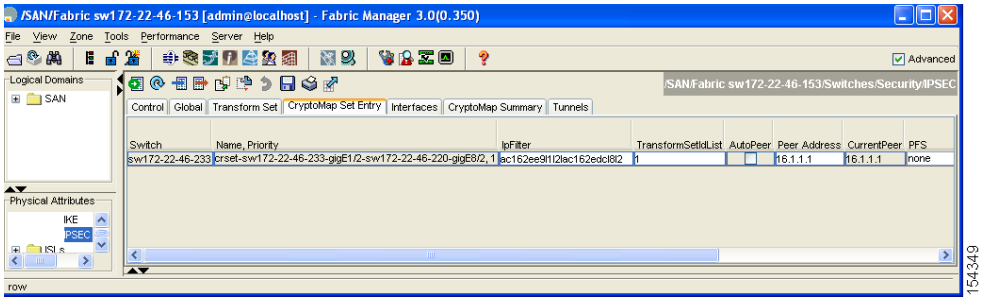

Step 4 Double-click and modify the value in the **Life Time(sec)** column.

Step 5 Click **Apply Changes** to save your changes.

About the AutoPeer Option

Setting the peer address as **AutoPeer** in the crypto map indicates that the destination endpoint of the traffic should be used as the peer address for the SA. Using the same crypto map, a unique SA can be set up to each of the endpoints in the subnet specified by the crypto map's IPv4 ACL entry. Auto-peer simplifies configuration when traffic endpoints are IPsec capable. It is particularly useful for iSCSI, where the iSCSI hosts in the same subnet do not require separate configuration.

Figure 39-27 shows a scenario where the auto-peer option can simplify configuration. Using the auto-peer option, only one crypto map entry is needed for all the hosts from subnet X to set up SAs with the switch. Each host will set up its own SA, but will share the crypto map entry. Without the auto-peer option, each host needs one crypto map entry.

Figure 39-27 iSCSI with End-to-End IPsec Using the auto-peer Option

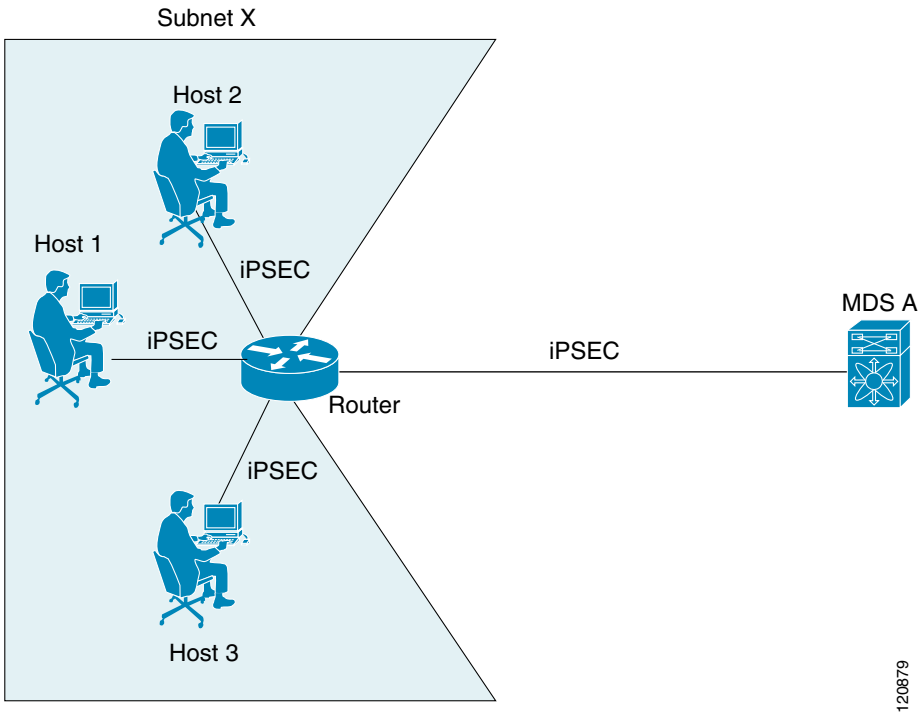

Configuring the AutoPeer Option

To configure the AutoPeer option using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **IPSEC** in the Physical Attributes pane.

You see the IPsec configuration in the Information pane.

Figure 39-28 IPsec Configuration

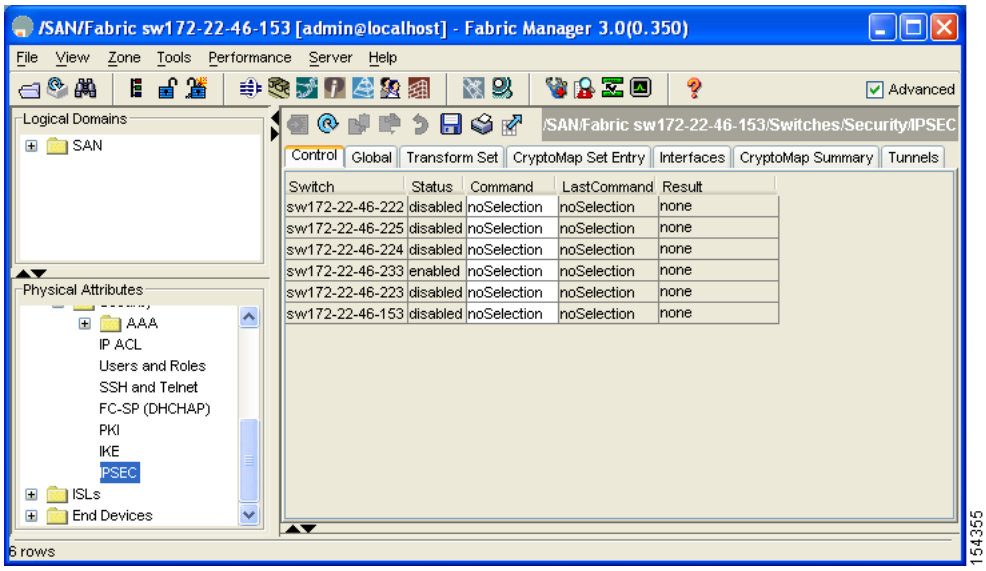

Step 2 Choose the **CryptoMap Set Entry** tab.

You see the existing crypto maps configured in Figure 39-29.

Figure 39-29 Existing Crypto Maps

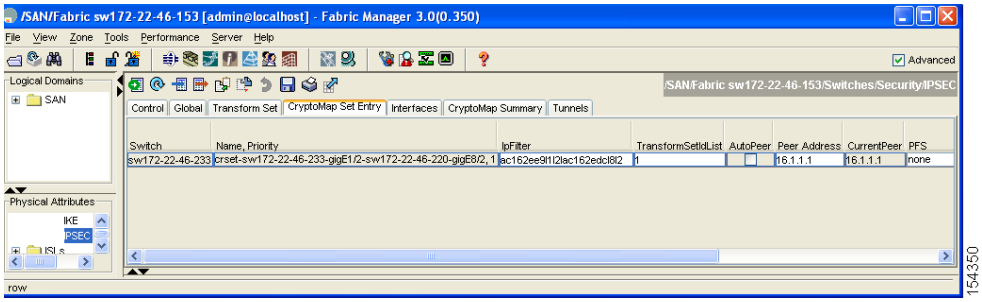

Step 3 Check or uncheck the **AutoPeer** option for the selected crypto map set entry.

Step 4 Click **Apply Changes** to save your changes.

About Perfect Forward Secrecy

To specify SA lifetime negotiation values, you can also optionally configure the perfect forward secrecy (PFS) value in the crypto map.

The PFS feature is disabled by default. If you set the PFS group, you can set one of the DH groups: 1, 2, 5, or 14. If you do not specify a DH group, the software uses group 1 by default.

Configuring Perfect Forward Secrecy

To configure the PFS value using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **IPSEC** in the Physical Attributes pane.

You see the IPsec configuration in the Information pane.

Figure 39-30 IPsec Configuration

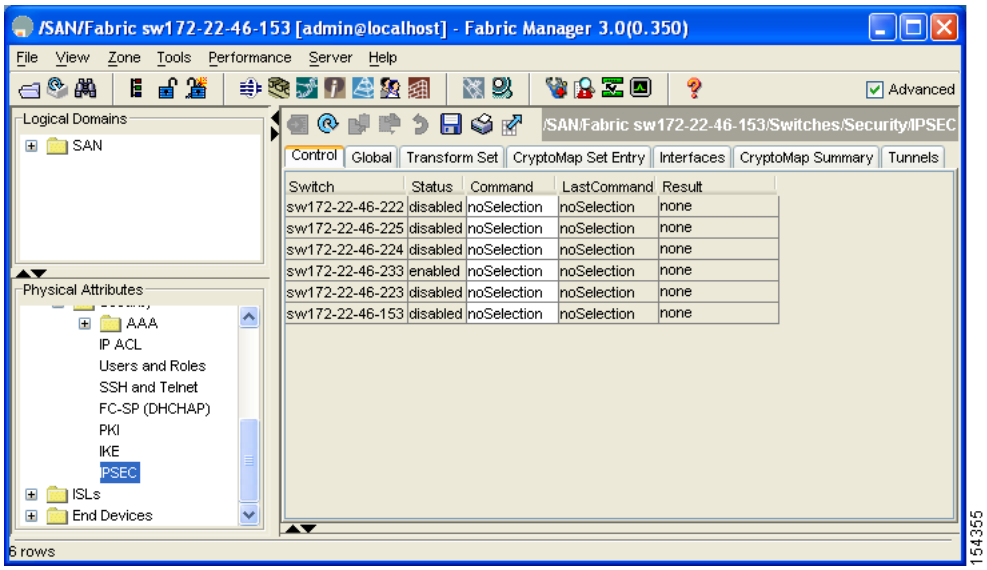

Step 2 Choose the **CryptoMap Set Entry** tab.

You see the existing crypto maps configured in Figure 39-31.

Figure 39-31 Existing Crypto Maps

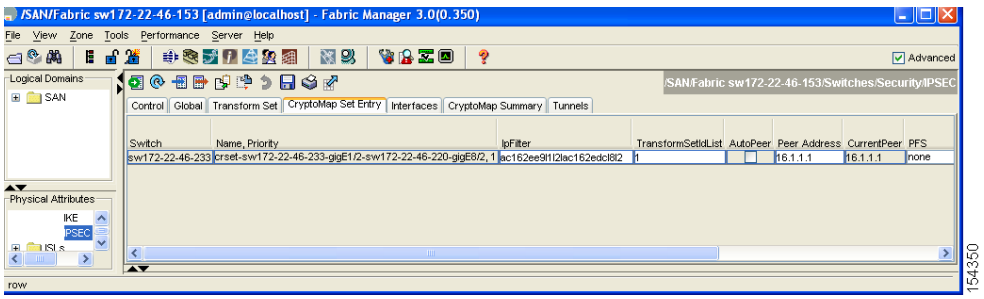

Step 3 Click the drop-down list in the **PFS** column and select the appropriate value.

Step 4 Click **Apply Changes** to save your changes.

Crypto Map Set Application

You need to apply a crypto map set to each interface through which IPsec traffic will flow. Applying the crypto map set to an interface instructs the switch to evaluate all the interface's traffic against the crypto map set and to use the specified policy during connection or SA negotiation on behalf of the traffic to be protected by crypto.

You can apply only one crypto map set to an interface. You can apply the same crypto map to multiple interfaces. However, you cannot apply more than one crypto map set to each interface.

Applying a Crypto Map Set

To apply a crypto map set to an interface using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **IPSEC** in the Physical Attributes pane.

You see the IPsec configuration in the Information pane in Figure 39-32.

Figure 39-32 IPsec Configuration

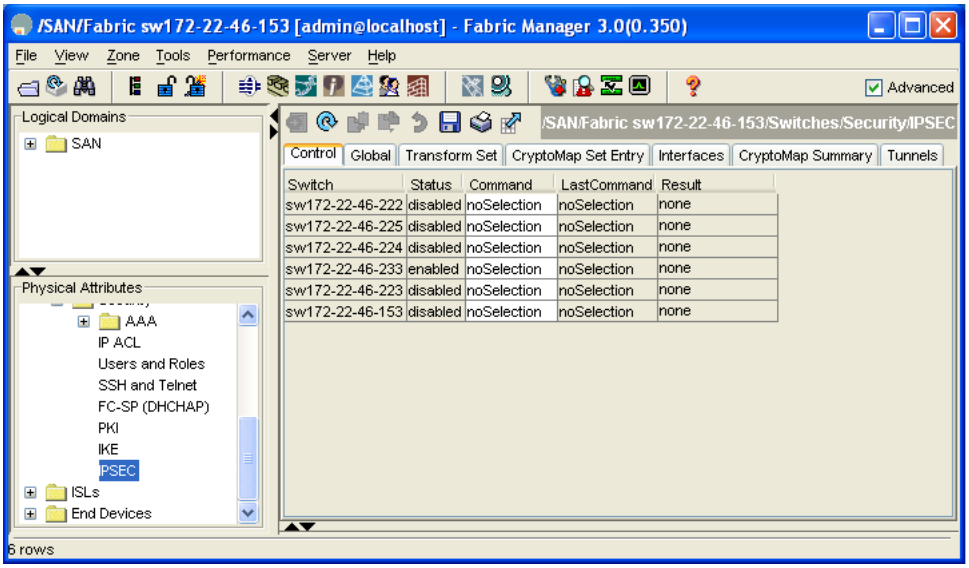

Step 2 Click the **Interfaces** tab.

You see the existing interface to crypto map configuration in Figure 39-33.

Figure 39-33 Crypto Map Interfaces

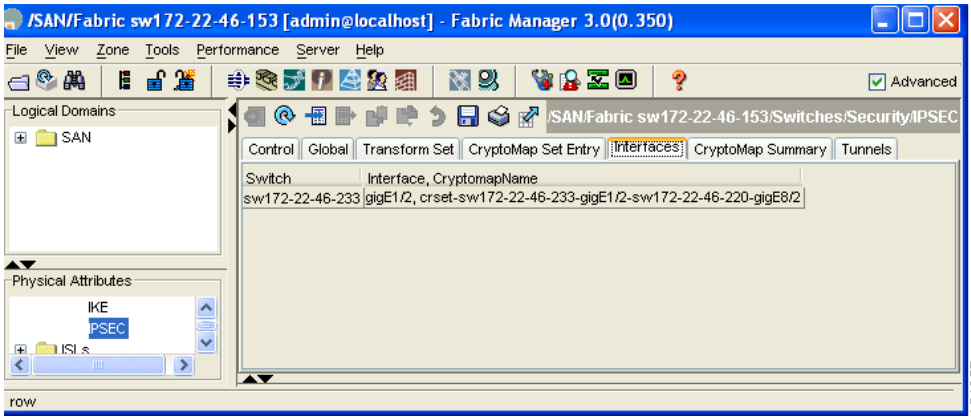

Step 3 Select the switch and interface you want to configure (see Figure 39-33).

Step 4 Enter the name of the crypto map that you want to apply to this interface in the CryptomapSetName field (see Figure 39-33).

Step 5 Click **Create** to apply the crypto map to the selected interface or click **Close** to exit the dialog box without applying the crypto map (see Figure 39-33).

IPsec Maintenance

Certain configuration changes will only take effect when negotiating subsequent security associations. If you want the new settings to take immediate effect, you must clear the existing security associations so that they will be reestablished with the changed configuration. If the switch is actively processing IPsec traffic, it is desirable to clear only the portion of the security association database that would be affected by the configuration changes (that is, clear only the security associations established by a given crypto map set). Clearing the full security association database should be reserved for large-scale changes, or when the router is processing very little other IPsec traffic.

Global Lifetime Values

If you have not configured a lifetime in the crypto map entry, the global lifetime values are used when negotiating new IPsec SAs.

You can configure two lifetimes: timed or traffic-volume. A SA expires after the first of these lifetimes is reached. The default lifetimes are 3,600 seconds (one hour) and 450 GB.

If you change a global lifetime, the new lifetime value will not be applied to currently existing SAs, but will be used in the negotiation of subsequently established SAs. If you wish to use the new values immediately, you can clear all or part of the SA database.

Assuming that the particular crypto map entry does not have lifetime values configured, when the switch requests new SAs it will specify its global lifetime values in the request to the peer; it will use this value as the lifetime of the new SAs. When the switch receives a negotiation request from the peer, it uses the value determined by the IKE version in use:

- If you use IKEv1 to setup IPsec SAs, the SA lifetime values are chosen to be the smaller of the two proposals. The same values are programmed on both the ends of the tunnel.
- If you use IKEv2 to setup IPsec SAs, SAs on each end has its own set up of lifetime values and thus the SAs on both sides expire independently.

The SA (and corresponding keys) will expire according to whichever comes sooner, either after the specified amount of time (in seconds) has passed or after the specified amount of traffic (in bytes) has passed.

A new SA is negotiated before the lifetime threshold of the existing SA is reached, to ensure that negotiation completes before the existing SA expires.

The new SA is negotiated when one of the following thresholds is reached (whichever comes first):

- 30 seconds before the lifetime expires or
- Approximately 10% of the lifetime in bytes remain

If no traffic has passed through when the lifetime expires, a new SA is not negotiated. Instead, a new SA will be negotiated only when IPsec sees another packet that should be protected.

To configure global SA lifetimes using Fabric Manager, follow these steps:

Step 1 Choose **Switches > Security** and then select **IPSEC** in the Physical Attributes pane.

You see the IP Sec configuration in the Information pane.

Step 2 Choose the **Global** tab.

Step 3 Double-click and edit the value in the **Life Time(sec)** column highlighted in Figure 39-34.

Figure 39-34 IP Sec Configuration Global Tab

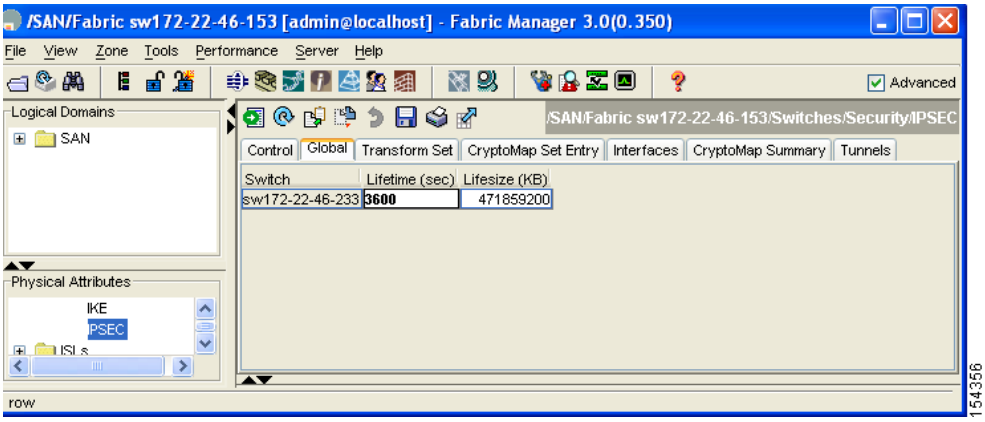

Step 4 Click **Apply Changes** to save your changes.

Default Settings

Table 39-3 lists the default settings for IKE parameters.

Table 39-3 Default IKE Parameters

Parameters	Default
IKE	Disabled.
IKE version	IKE version 2.
IKE encryption algorithm	3DES.
IKE hash algorithm	SHA.
IKE authentication method	Not configurable (uses preshared keys).
IKE DH group identifier	Group 1.
IKE lifetime association	86,400 00 seconds (equals 24 hours).
IKE keepalive time for each peer (v2)	3,600 seconds (equals one hour).

Table 39-4 lists the default settings for IPsec parameters.

Table 39-4 Default IPsec Parameters

Parameters	Default
IPsec	Disabled.
Applying IPsec to the traffic.	Deny—allowing clear text.
IPsec PFS	Disabled.
IPsec global lifetime (traffic-volume)	450 Gigabytes.
IPsec global lifetime (time)	3,600 seconds (one hour).

CHAPTER 40

Configuring FC-SP and DHCHAP

Fibre Channel Security Protocol (FC-SP) capabilities provide switch-switch and host-switch authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) is an FC-SP protocol that provides authentication between Cisco MDS 9000 Family switches and other devices. DHCHAP consists of the CHAP protocol combined with the Diffie-Hellman exchange.

This chapter includes the following sections:

About Fabric Authentication

All switches in the Cisco MDS 9000 Family enable fabric-wide authentication from one switch to another switch, or from a switch to a host. These switch and host authentications are performed locally or remotely in each fabric. As storage islands are consolidated and migrated to enterprise-wide fabrics new security challenges arise. The approach of securing storage islands cannot always be guaranteed in enterprise-wide fabrics. For example, in a campus environment with geographically distributed switches someone could maliciously interconnect incompatible switches or you could accidentally do so, resulting in Inter-Switch Link (ISL) isolation and link disruption. This need for physical security is addressed by switches in the Cisco MDS 9000 Family (see Figure 40-1).

Figure 40-1 Switch and Host Authentication

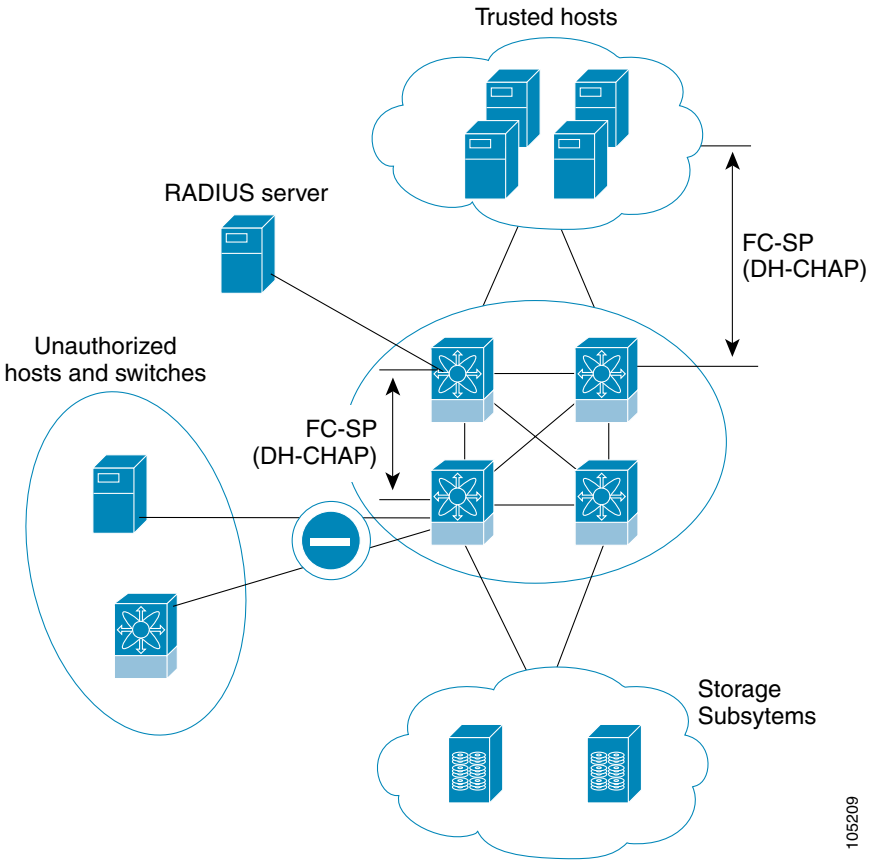

Note Fibre Channel (FC) host bus adapters (HBAs) with appropriate firmware and drivers are required for host-switch authentication.

DHCHAP

DHCHAP is an authentication protocol that authenticates the devices connecting to a switch. Fibre Channel authentication allows only trusted devices to be added to a fabric, thus preventing unauthorized devices from accessing the switch.

Note The terms FC-SP and DHCHAP are used interchangeably in this chapter.

DHCHAP is a mandatory password-based, key-exchange authentication protocol that supports both switch-to-switch and host-to-switch authentication. DHCHAP negotiates hash algorithms and DH groups before performing authentication. It supports MD5 and SHA-1 algorithm-based authentication.

Configuring the DHCHAP feature requires the ENTERPRISE_PKG license (see Chapter 10, "Obtaining and Installing Licenses").

This sections identifies the impact of configuring the DHCHAP feature along with existing Cisco MDS features:

- PortChannel interfaces—If DHCHAP is enabled for ports belonging to a PortChannel, DHCHAP authentication is performed at the physical interface level, not at the PortChannel level.
- FCIP interfaces—The DHCHAP protocol works with the FCIP interface just as it would with a physical interface.
- Port security or fabric binding—Fabric binding policies are enforced based on identities authenticated by DHCHAP.
- VSANs—DHCHAP authentication is not done on a per-VSAN basis.
- High availability—DHCHAP authentication works transparently with existing HA features.

To configure DHCHAP authentication using the local password database, follow these steps:

Step 1 Enable DHCHAP.

Step 2 Identify and configure the DHCHAP authentication modes.

Step 3 Configure the hash algorithm and DH group.

Step 4 Configure the DHCHAP password for the local switch and other switches in the fabric.

Step 5 Configure the DHCHAP time out value for reauthentication.

Step 6 Verify the DHCHAP configuration.

This section contains the following topics:

About Enabling DHCHAP

By default, the DHCHAP feature is disabled in all switches in the Cisco MDS 9000 Family.

You must explicitly enable the DHCHAP feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.

Enabling DHCHAP

To enable DHCHAP for a Cisco MDS switch using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **FC-SP**.

You see the FC-SP (DHCHAP) configuration in the Information pane shown in Figure 40-2.

Figure 40-2 FC-SP (DHCHAP) Configuration

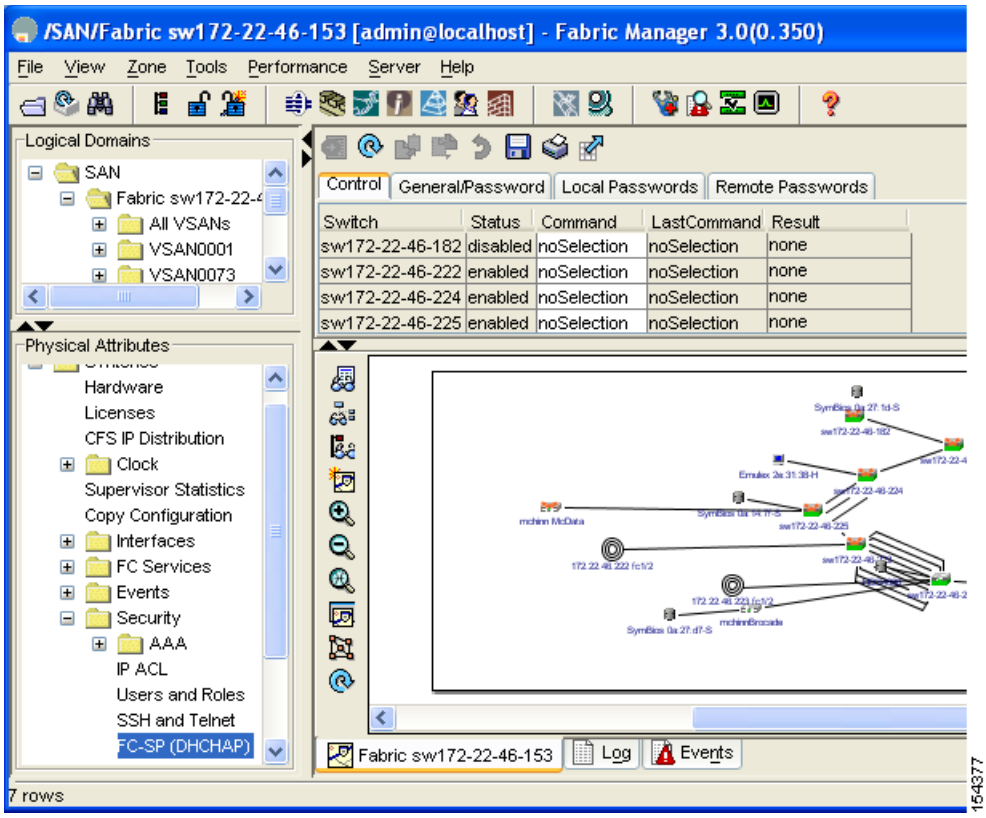

The **Control** tab is default. You see the FC-SP enable state for all switches in the fabric.

Step 2 Set the Command drop-down menu to **enable** for all switches that you want to enable FC-SP on.

Step 3 Click **Apply Changes** to enable FC-SP and DHCHAP on the selected switches.

About DHCHAP Authentication Modes

The DHCHAP authentication status for each interface depends on the configured DHCHAP port mode.

When the DHCHAP feature is enabled in a switch, each Fibre Channel interface or FCIP interface may be configured to be in one of four DHCHAP port modes:

- On—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software moves the link to an isolated state.
- Auto-Active—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software continues with the rest of the initialization sequence.
- Auto-Passive (default)—The switch does not initiate DHCHAP authentication, but participates in DHCHAP authentication if the connecting device initiates DHCHAP authentication.
- Off—The switch does not support DHCHAP authentication. Authentication messages sent to such ports return error messages to the initiating switch.

Note Whenever DHCHAP port mode is changed to a mode other than the Off mode, reauthentication is performed.

Table 40-1 identifies the switch-to-switch authentication behavior between two Cisco MDS switches in various modes.

Table 40-1 DHCHAP Authentication Status Between Two MDS Switches

<table>
<tr><th rowspan="2">Switch N DHCHAP Modes</th><th colspan="4">Switch 1 DHCHAP Modes</th></tr>
<tr><th>on</th><th>auto-active</th><th>auto-passive</th><th>off</th></tr>
<tr><td>on</td><td rowspan="3">FC-SP authentication is performed.</td><td rowspan="3">FC-SP authentication is performed.</td><td rowspan="2">FC-SP authentication is performed.</td><td>Link is brought down.</td></tr>
<tr><td>auto-Active</td><td rowspan="2">FC-SP authentication is not performed.</td></tr>
<tr><td>auto-Passive</td><td></td></tr>
<tr><td>off</td><td>Link is brought down.</td><td colspan="3">FC-SP authentication is not performed.</td></tr>
</table>

Configuring the DHCHAP Mode

To configure the DHCHAP mode for a particular interface using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **FC-SP**.

You see the FC-SP (DHCHAP) configuration in the Information pane shown in Figure 40-2.

Figure 40-3 FC-SP Configuration

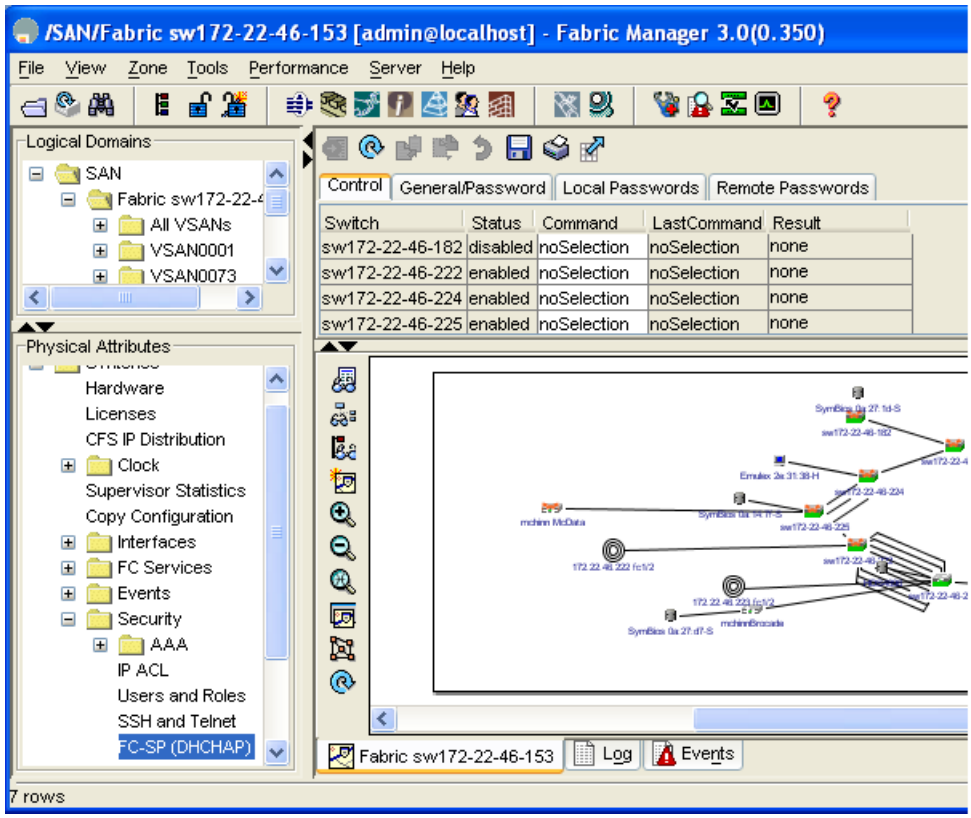

Step 2 Set the **Command** drop-down menu to the DHCHAP authentication mode you want to configure for that interface.

Step 3 Click **Apply Changes** to save these DHCHAP port mode settings.

About the DHCHAP Hash Algorithm

Cisco MDS switches support a default hash algorithm priority list of MD5 followed by SHA-1 for DHCHAP authentication.

Tip If you change the hash algorithm configuration, then change it globally for all switches in the fabric.

Caution RADIUS and TACACS+ protocols always use MD5 for CHAP authentication. Using SHA-1 as the hash algorithm may prevent RADIUS and TACACS+ usage—even if these AAA protocols are enabled for DHCHAP authentication.

Configuring the DHCHAP Hash Algorithm

To configure the hash algorithm using Fabric Manager, follow these steps:

Step 1 Choose **Switches > Security** and then select **FC-SP**.

You see the FC-SP configuration in the Information pane shown in Figure 40-4.

Figure 40-4 FC-SP Configuration

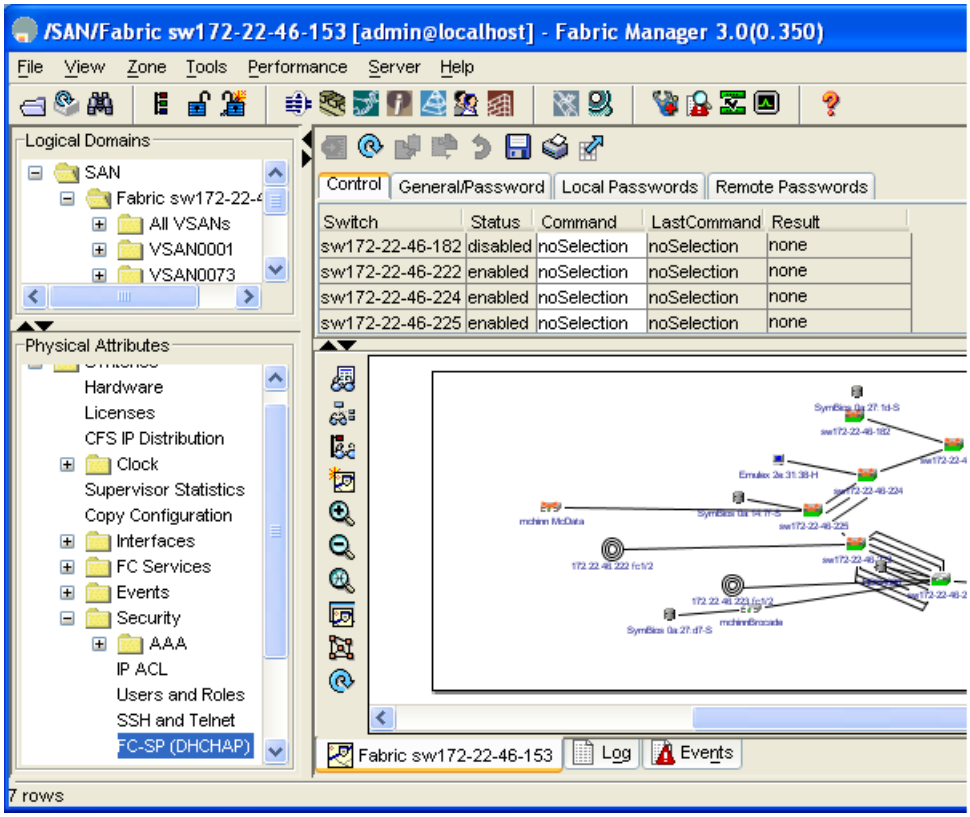

Step 2 Click the **General/Password** tab.

You see the DHCHAP general settings mode for each switch shown in Figure 40-5.

Figure 40-5 General/ Password Tab

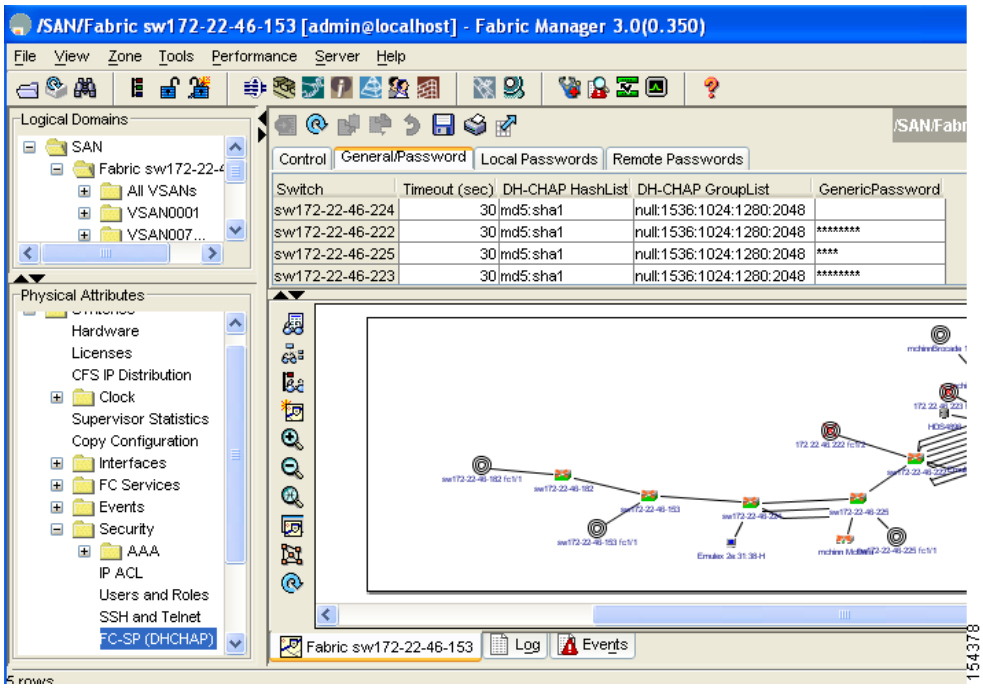

Step 3 Change the DHCHAP HashList (see Figure 40-5) for each switch in the fabric.

Step 4 Click **Apply Changes** to save the updated hash algorithm priority list or click **Undo Changes** to discard any unsaved changes.

About the DHCHAP Group Settings

All switches in the Cisco MDS Family support all DHCHAP groups specified in the standard: 0 (null DH group, which does not perform the Diffie-Hellman exchange), 1, 2, 3, or 4.

Tip If you change the DH group configuration, change it globally for all switches in the fabric.

Configuring the DHCHAP Group Settings

To change the DH group settings using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **FC-SP**.

You see the FC-SP configuration in the Information pane.

Figure 40-6 FC-SP Configuration

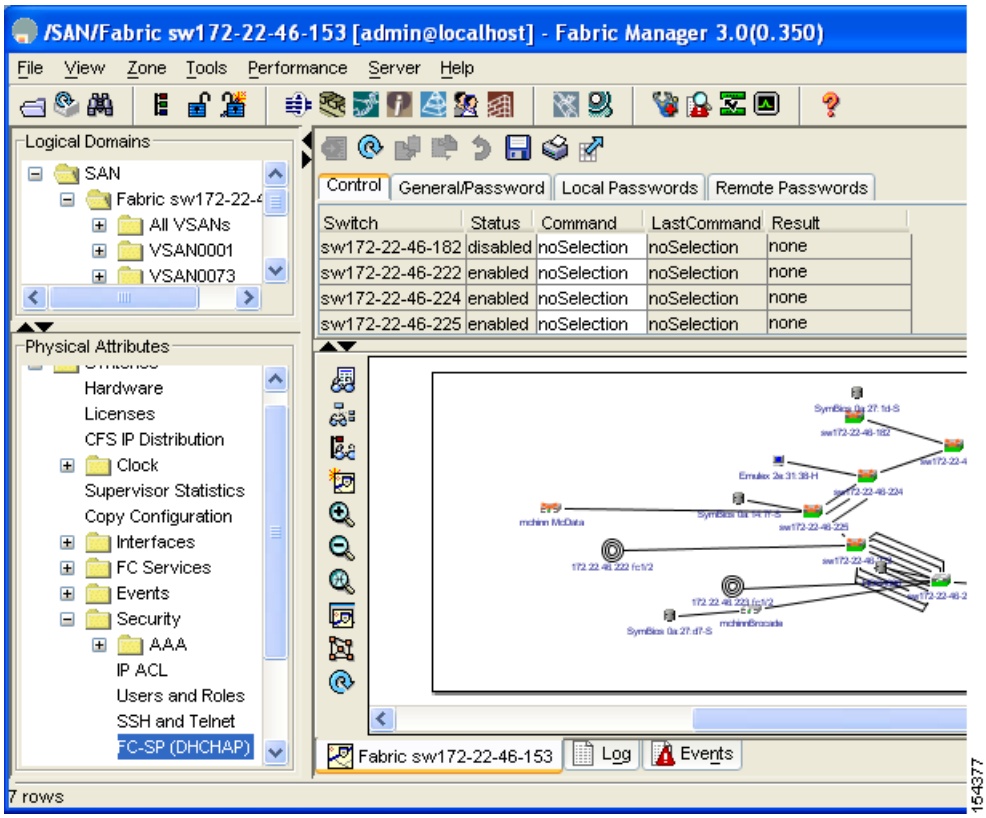

Step 2 Choose the **General/Password** tab.

You see the DHCHAP general settings mode for each switch.

Step 3 Change the DH CHAP GroupList for each switch in the fabric.

Step 4 Click **Apply Changes** icon to save the updated hash algorithm priority list or click **Undo Changes** to discard any unsaved changes.

About the DHCHAP Password Configuration

DHCHAP authentication in each direction requires a shared secret password between the connected devices. To do this, you can use one of three approaches to manage passwords for all switches in the fabric that participate in DHCHAP.

- Approach 1—Use the same password for all switches in the fabric. This is the simplest approach. When you add a new switch, you use the same password to authenticate that switch in this fabric. It is also the most vulnerable approach if someone from the outside maliciously attempts to access any one switch in the fabric.
- Approach 2—Use a different password for each switch and maintain that password list in each switch in the fabric. When you add a new switch, you create a new password list and update all switches with the new list. Accessing one switch yields the password list for all switches in that fabric.
- Approach 3—Use different passwords for different switches in the fabric. When you add a new switch, multiple new passwords corresponding to each switch in the fabric must be generated and configured in each switch. Even if one switch is compromised, the password of other switches are still protected. This approach requires considerable password maintenance by the user.

Note All passwords are restricted to 64 alphanumeric characters and can be changed, but not deleted.

Tip We recommend using RADIUS or TACACS+ for fabrics with more than five switches. If you need to use a local password database, you can continue to do so using Approach 3 and using the Cisco MDS 9000 Family Fabric Manager to manage the password database.

Configuring DHCHAP Passwords for the Local Switch

To configure the DHCHAP password for the local switch using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **FC-SP**.

You see the FC-SP configuration in the Information pane.

Step 2 Click the **Local Passwords** tab.

You see existing DHCHAP local passwords for each switch.

Step 3 Click the **Create Row** icon to create a new local password.

You see the Create Local Passwords dialog box.

Step 4 Optionally, check the switches that you want to configure the same local password on.

Step 5 Select the switch WNN and fill in the Password field.

Step 6 Click Create to save the updated password or click Close to discard any unsaved changes.

About Password Configuration for Remote Devices

You can configure passwords in the local authentication database for other devices in a fabric. The other devices are identified by their device name, which is also known as the switch WWN or device WWN. The password is restricted to 64 characters and can be specified in clear text (0) or in encrypted text (7).

Note The switch WWN identifies the physical switch. This WWN is used to authenticate the switch and is different from the VSAN node WWN.

Configuring DHCP Passwords for Remote Devices

To locally configure the remote password for another switch in the fabric using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **FC-SP**.

You see the FC-SP configuration in the Information pane.

Step 2 Click the **Remote Passwords** tab.

You see the DHCHAP local password for each switch shown in Figure 40-7.

Figure 40-7 Remote Passwords for DHCHAP

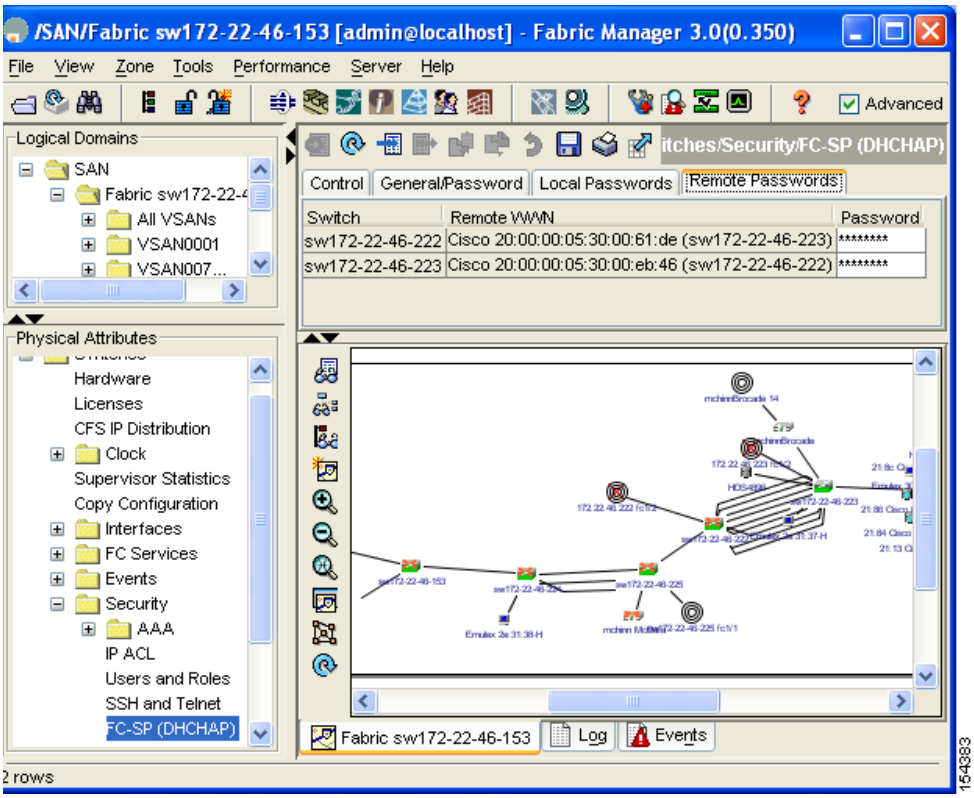

Step 3 Click **Create Row** to create a remote password.

You see the Create Remote Passwords dialog box.

Step 4 Optionally, check the switches that you want to configure the same remote password on in Fabric Manager.

Step 5 Select the switch WNN for the remote device and fill in the Password field.

Step 6 Click **Create** to save the updated password or click **Close** to discard any unsaved changes.

About the DHCHAP Time Out Value

During the DHCHAP protocol exchange, if the MDS switch does not receive the expected DHCHAP message within a specified time interval, authentication failure is assumed. The time ranges from 20 (no authentication is performed) to 1000 seconds. The default is 30 seconds.

When changing the time out value, consider the following factors:

- The existing RADIUS and TACACS+ time out values.
- The same value must also be configured on all switches in the fabric.

Configuring the DHCHAP Time Out Value

To configure the DHCHAP time out value using Fabric Manager, follow these steps:

Step 1 Expand **Switches > Security** and then select **FC-SP**.

You see the FC-SP configuration in the Information pane.

Step 2 Choose the **General/Password** tab.

You see the DHCHAP general settings mode for each switch.

Step 3 Change the DHCHAP time out value for each switch in the fabric.

Step 4 Click **Apply Changes** to save the updated hash algorithm priority list or click **Undo Changes** to discard any unsaved changes.

Configuring DHCHAP AAA Authentication

You can individually set authentication options. If authentication is not configured, local authentication is used by default.

To configure the AAA authentication refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

Enabling FC-SP on ISLs

There is anISL pop-up menu in Fabric Manager called Enable FC-SP that enables FC-SP on switches at either end of the ISL. You are prompted for an FC-SP generic password, then asked to set FC-SP interface mode to ON for affected ports. Right-click an ISL and click **Enable FC-SP** to access this feature.

Default Settings

Table 40-2 lists the default settings for all fabric security features in any switch.

Table 40-2 Default Fabric Security Settings

Parameters	Default
DHCHAP feature	Disabled.
DHCHAP hash algorithm	A priority list of MD5 followed by SHA-1 for DHCHAP authentication.
DHCHAP authentication mode	Auto-passive.
DHCHAP group default priority exchange order	0, 4, 1, 2, and 3 respectively.
DHCHAP time out value	30 seconds.

CHAPTER 41

Configuring Port Security

All switches in the Cisco MDS 9000 Family provide port security features that reject intrusion attempts and report these intrusions to the administrator.

Note Port security is only supported for Fibre Channel ports.

This chapter includes the following sections:

About Port Security

Typically, any Fibre Channel device in a SAN can attach to any SAN switch port and access SAN services based on zone membership. Port security features prevent unauthorized access to a switch port in the Cisco MDS 9000 Family:

- Login requests from unauthorized Fibre Channel devices (Nx ports) and switches (xE ports) are rejected.
- All intrusion attempts are reported to the SAN administrator through system messages.
- Configuration distribution uses the CFS infrastructure, and is limited to those switches that are CFS capable. Distribution is disabled by default.
- Configuring the port security policy requires the ENTERPRISE_PKG license (see Chapter 10, “Obtaining and Installing Licenses”).

This section includes the following topics:

Port Security Enforcement

To enforce port security, configure the devices and switch port interfaces through which each device or switch is connected, and activate the configuration.

- Use the port world wide name (pWWN) or the node world wide name (nWWN) to specify the Nx port connection for each device.
- Use the switch world wide name (sWWN) to specify the xE port connection for each switch.

Each Nx and xE port can be configured to restrict a single port or a range of ports.

Enforcement of port security policies are done on every activation and when the port tries to come up.

The port security feature uses two databases to accept and implement configuration changes.

- Configuration database—All configuration changes are stored in the configuration database.
- Active database—The database currently enforced by the fabric. The port security feature requires all devices connecting to a switch to be part of the port security active database. The software uses this active database to enforce authorization.

About Auto-Learning

You can instruct the switch to automatically learn (auto-learn) the port security configurations over a specified period. This feature allows any switch in the Cisco MDS 9000 Family to automatically learn about devices and switches that connect to it. Use this feature when you activate the port security feature for the first time as it saves tedious manual configuration for each port. You must configure auto-learning on a per-VSAN basis. If enabled, devices and switches that are allowed to connect to the switch are automatically learned, even if you have not configured any port access.

When auto-learning is enabled, learning happens only for the devices or interfaces that were not already logged into the switch. Learned entries on a port are cleaned up after you shut down that port if auto-learning is still enabled.

Learning does not override the existing configured port security policies. So, for example, if an interface is configured to allow a specific pWWN, then auto-learning will not add a new entry to allow any other pWWN on that interface. All other pWWNs will be blocked even in auto-learning mode.

No entries are learned for a port in the shutdown state.

When you activate the port security feature, auto-learning is also automatically enabled.

Note If you enable auto-learning before activating port security, you cannot activate until auto-learning is disabled.

Port Security Activation

By default, the port security feature is not activated in any switch in the Cisco MDS 9000 Family.

By activating the port security feature, the following apply:

- Auto-learning is also automatically enabled, which means:
  - From this point, auto-learning happens only for the devices or interfaces that were not logged into the switch.
  - You cannot activate the database until you disable auto-learning.
- All the devices that are already logged in are learned and are added to the active database.
- All entries in the configured database are copied to the active database.

After the database is activated, subsequent device login is subject to the activated port bound WWN pairs, excluding the auto-learned entries. You must disable auto-learning before the auto-learned entries become activated.

When you activate the port security feature, auto-learning is also automatically enabled. You can choose to activate the port security feature and disable auto-learning.

Tip If a port is shut down because of a denied login attempt, and you subsequently configure the database to allow that login, the port does not come up automatically. You must explicitly issue a **no shutdown** CLI command to bring that port back online.

Port Security Configuration Guidelines

To enforce port security, configure the devices and switch port interfaces through which each device or switch is connected, and activate the configuration.

- Use the port world wide name (pWWN) or the node world wide name (nWWN) to specify the Nx port connection for each device.
- Use the switch world wide name (sWWN) to specify the xE port connection for each switch.

Each Nx and xE port can be configured to restrict a single port or a range of ports.

Enforcement of port security policies are done on every activation and when the port tries to come up.

The port security feature uses two databases to accept and implement configurations.

- Configuration database—All configuration changes are stored in the configuration database.
- Active database—The database currently enforced by the fabric. The port security feature requires all devices connecting to a switch to be part of the port security active database. The software uses this active database to enforce authorization.

The steps to configure port security depend on which features you are using. Auto-learning works differently if you are using CFS distribution.

This section includes the following topics:

Configuring Port Security with Auto-Learning and CFS Distribution

To configure port security, using auto-learning and CFS distribution, follow these steps:

Step 1 Enable port security. See the "Enabling Port Security" section on page 41-5.

Step 2 Enable CFS distribution. See the "Enabling Distribution" section on page 41-17.

Step 3 Activate port security on each VSAN. This turns on auto-learning by default. See the "Activating Port Security" section on page 41-7.

Step 4 Issue a CFS commit to copy this configuration to all switches in the fabric. See the "Committing the Changes" section on page 41-18. At this point, all switches are activated, and auto-learning.

Step 5 Wait until all switches and all hosts are automatically learned.

Step 6 Disable auto-learn on each VSAN. See the "Disabling Auto-learning" section on page 41-12.

Step 7 Issue a CFS commit to copy this configuration to all switches in the fabric. See the "Committing the Changes" section on page 41-18. At this point, the auto-learned entries from every switch are combined into a static active database that is distributed to all switches.

Step 8 Copy the active database to the configure database on each VSAN. See the "Port Security Database Copy" section on page 41-22.

Step 9 Issue a CFS commit to copy this configuration to all switches in the fabric. See the "Committing the Changes" section on page 41-18. This ensures that the configure database is the same on all switches in the fabric.

Step 10 Copy the running configuration to the startup configuration, using the fabric option. This saves the port security configure database to the startup configuration on all switches in the fabric.

Configuring Port Security with Auto-Learning without CFS

To configure port security using auto-learning without CFS, follow these steps:

Step 1 Enable port security. See the "Enabling Port Security" section on page 41-5.

Step 2 Activate port security on each VSAN. This turns on auto-learning by default. See the "Activating Port Security" section on page 41-7.

Step 3 Wait until all switches and all hosts are automatically learned.

Step 4 Disable auto-learn on each VSAN. See the "Disabling Auto-learning" section on page 41-12.

Step 5 Copy the active database to the configure database on each VSAN. See the "Port Security Database Copy" section on page 41-22.

Step 6 Copy the running configuration to the startup configuration This saves the port security configure database to the startup configuration.

Step 7 Repeat Step 1 through Step 6 for all switches in the fabric.

Configuring Port Security with Manual Database Configuration

To configure port security and manually configure the port security database, follow these steps:

Step 1 Enable port security. See the "Enabling Port Security" section on page 41-5.

Step 2 Manually configure all port security entries into the configure database on each VSAN. See the "Configuring Port Security with Manual Database Configuration" section on page 41-5.

Step 3 Activate port security on each VSAN. This turns on auto-learning by default. See the "Activating Port Security" section on page 41-7.

Step 4 Disable auto-learn on each VSAN. See the "Disabling Auto-learning" section on page 41-12.

Step 5 Copy the running configuration to the startup configuration This saves the port security configure database to the startup configuration.

Step 6 Repeat Step 1 through Step 5 for all switches in the fabric.

Enabling Port Security

By default, the port security feature is disabled in all switches in the Cisco MDS 9000 Family.

To enable port security using Fabric Manager, follow these steps:

Step 1 Expand a **VSAN** and then select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane. See Figure 41-1.

Figure 41-1 Port Security Configuration

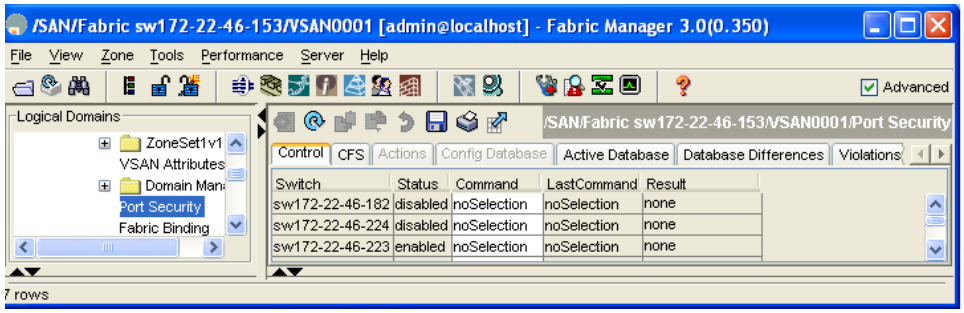

Step 2 Click the **CFS** tab.

You see the information in Figure 41-2.

Figure 41-2 Port Security CFS

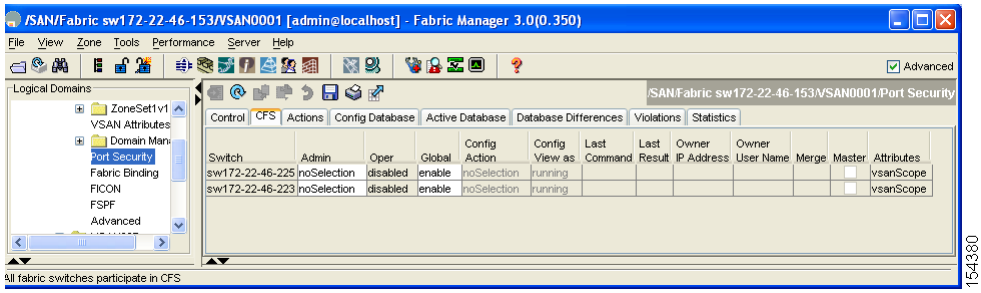

Step 3 Enable CFS on all participating switches in the VSAN by clicking each entry in the Global column and selecting **enable**.

Step 4 Click **Apply Changes** to enable CFS distribution for the port security feature.

Step 5 Click the **Control** tab.

You see the port security enable state for all switches in the selected VSAN (see Figure 41-3).

Figure 41-3 Port Security Configuration

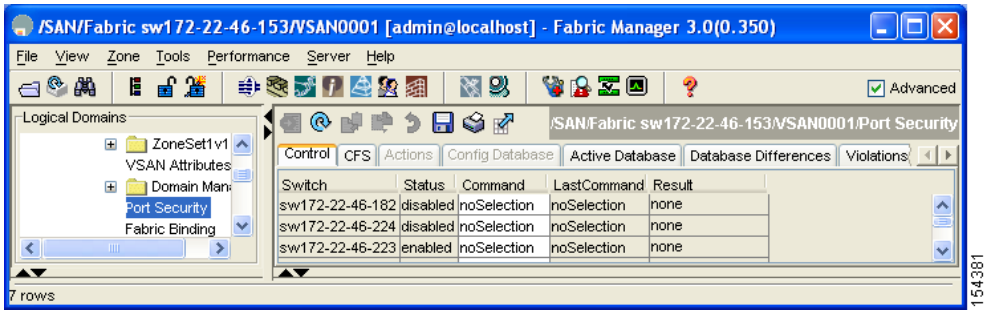

Step 6 Set the Command column (see Figure 41-3) to **enable** for each switch in the VSAN.

Step 7 Click the **CFS** tab and set the Command column to **commit** on all participating switches in the VSAN.

Step 8 Click **Apply Changes** to distribute the enabled port security to all switches in the VSAN.

About Port Security Activation

By default, the port security feature is not activated in any switch in the Cisco MDS 9000 Family.

By activating the port security feature, the following apply:

- Auto-learning is also automatically enabled, which means:
  - From this point, auto-learning happens only for the devices or interfaces that were not activated.
  - You cannot activate the database until you disable auto-learning.
- All the logged-in devices are learned and are added to the active database.
- All entries in the configured database are copied to the active database.

After the database is activated, subsequent device login is subject to the activated port bound WWN pairs.

When you activate the port security feature, auto-learning is also automatically enabled. You can choose to activate the port security feature and disable auto-learning.

This section includes the following topics:

Activating Port Security

To activate port security using Fabric Manager, follow these steps:

Step 1 Expand a **VSAN** and select **Port Security** in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Actions** tab.

Step 3 Click in the Action column under Activation, next to the switch or VSAN on which you want to activate port security. You see a drop-down menu with the following options:

- **activate**—Valid port security settings are activated.
- **activate (TurnLearningOff)**—Valid port security settings are activated and auto-learn turned off.
- **forceActivate**—Activation is forced.
- **forceActivate(TurnLearningOff)**—Activation is forced and auto-learn is turned off.
- **deactivate**—All currently active port security settings are deactivated.
- **NoSelection**— No action is taken.

Step 4 Set the Action field you want for that switch.

Step 5 Uncheck the **AutoLearn** check box for each switch in the VSAN to disable auto-learning.

Step 6 Click the **CFS** tab and set the command column to **commit** on all participating switches in the VSAN.

Step 7 Click **Apply Changes** in Fabric Manager or **Apply** in Device Manager to save these changes or click **Undo Changes** in Fabric Manager or **Close** in Device Manager to discard any unsaved changes.

Note If required, you can disable auto-learning (see the "Disabling Auto-learning" section on page 41-12).

Database Activation Rejection

Database activation is rejected in the following cases:

- Missing or conflicting entries exist in the configuration database but not in the active database.
- If the auto-learning feature was enabled before the activation.
- The exact security is not configured for each PortChannel member.
- The configured database is empty but the active database is not.

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed by forcing the port security activation.

Forcing Port Security Activation

If the port security activation request is rejected, you can force the activation.

Note An activation using the **force** option can log out existing devices if they violate the active database.

To forcefully activate the port security database, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Actions** tab.

Step 3 Click in the **Action** column under Activation, next to the switch or VSAN on which you want to activate port security and select the **forceactivate** option.

Step 4 Set the Action field you want for that switch.

Step 5 Click the **CFS** tab and set the command column to **commit** on all participating switches in the VSAN.

Step 6 Click **Apply Changes** in Fabric Manager or **Apply** in Device Manager to save these changes or click **Undo Changes** in Fabric Manager or **Close** in Device Manager to discard any unsaved changes.

Database Reactivation

Tip If auto-learning is enabled and you activate the database, you will not be allowed to proceed.

To reactivate the database using Fabric Manager, follow these steps:

Step 1 Disable auto-learning.

Step 2 Copy the active database to the configured database.

Tip If the active database is empty, you cannot perform this step.

Step 3 Make the required changes to the configuration database.

Step 4 Activate the database.

Copying an Active Database to the Config Database

To copy the active database to the config database using Fabric Manager, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Actions** tab.

You see the switches for that VSAN.

Step 3 Check the **CopyActive ToConfig** check box next to the switch for which you want to copy the database.

The active database is copied to the config database when the security setting is activated.

Step 4 Uncheck the **CopyActive ToConfig** check box if you do not want the database copied when the security setting is activated.

Step 5 Click the **CFS** tab and set the command column to **commit** on all participating switches in the VSAN.

Step 6 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

Displaying Activated Port Security Settings

To display active port security settings using Fabric Manager, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Active Database** tab.

You see the active port security settings for that VSAN.

Displaying Port Security Statistics

To display port security statistics, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Statistics** tab.

You see the port security statistics for that VSAN.

Displaying Port Security Violations

Port violations are invalid login attempts (for example, login requests from unauthorized Fibre Channel devices). You can display a list of these attempts on a per-VSAN basis, using Fabric Manager.

To display port security violations, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Violations** tab. You see the port security violations for that VSAN.

About Auto-learning

You can instruct the switch to automatically learn (auto-learn) the port security configurations over a specified period. This feature allows any switch in the Cisco MDS 9000 Family to automatically learn about devices and switches that connect to it. Use this feature to activate the port security feature for the first time as it saves tedious manual configuration for each port. You must configure auto-learning on a per-VSAN basis. If enabled, devices and switches that are allowed to connect to the switch are automatically learned, even if you have not configured any port access. Learned entries on a port are cleaned up after you shut down that port. Learning does not override the enforced port security policies.

When you activate the port security feature, auto-learning is also automatically enabled. When auto-learning is enabled, the following apply:

- Learning happens only for the devices or interfaces that were not activated.
- You cannot activate the database.

This section contains the following topics:

About Enabling Auto-learning

The state of the auto-learning configuration depends on the state of the port security feature:

- If the port security feature is not activated, auto-learning is disabled by default.
- If the port security feature is activated, auto-learning is enabled by default (unless you explicitly disabled this option).

Tip If auto-learning is enabled on a VSAN, you can only activate the database for that VSAN by using the **force** option.

Enabling Auto-learning

To enable auto-learning using Fabric Manager, follow these steps:

Step 1 Expand a **VSAN** and select **Port Security** in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Figure 41-4 Port Security Configuration

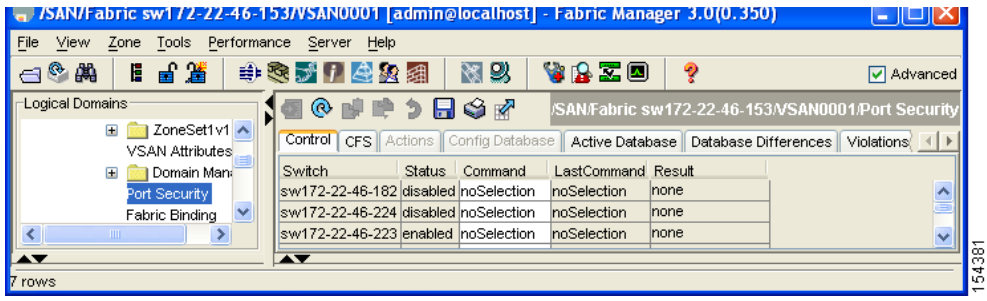

Step 2 Click the **Actions** tab.

Step 3 Click in the Action column under Activation, next to the switch or VSAN on which you want to activate port security. You see a drop-down menu with the following options:

- **activate**—Valid port security settings are activated.
- **activate (TurnLearningOff)**—Valid port security settings are activated and auto-learn turned off.
- **forceActivate**—Activation is forced.
- **forceActivate(TurnLearningOff)**—Activation is forced and auto-learn is turned off.
- **deactivate**—All currently active port security settings are deactivated.
- **NoSelection**— No action is taken.

Step 4 Select one of the port security options for that switch.

Step 5 Check the **AutoLearn** check box for each switch in the VSAN to enable auto-learning. Click the **CFS** tab and set the command column to **commit** on all participating switches in the VSAN.

Step 6 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

Disabling Auto-learning

To disable auto-learning using Fabric Manager, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Figure 41-5 Port Security Configuration

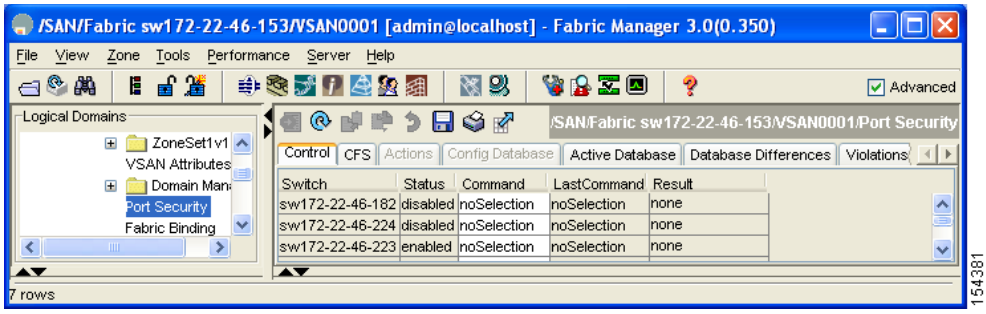

Step 2 Click the **Action** tab.

You see the switches for that VSAN.

Step 3 Check the **AutoLearn** check box next to the switch if you want to enable auto-learning.

Step 4 Uncheck the **AutoLearn** check box next to the switch if you want to disable auto-learning.

Step 5 Click the **CFS** button at the top of the Information pane and select **commit**.

Step 6 Click the **Apply Changes** icon to save these changes or click **Undo Changes** to discard any unsaved changes.

Auto-learning Device Authorization

Table 41-1 summarizes the authorized connection for device requests.

Table 41-1 Authorized Auto-learning Device Requests

Device (pWWN, nWWN, sWWN)	Requests Connection to	Authorization	Condition
Configured with one or more switch ports	A switch on configured ports	Permitted	1
	A switch on other ports	Denied	2
Not configured	A port that is not configured	Permitted if auto-learning enabled	3
		Denied if auto-learning disabled	4
Configured or not configured	A switch port that allows any device	Permitted	5
Configured to log in to any switch port	Any port on the switch	Permitted	6
Not configured	A port configured with some other device	Denied	7

Authorization Scenario

Assume that the port security feature is activated and the following conditions are specified in the active database:

- A pWWN (P1) is allowed access through interface fc1/1 (F1).
- A pWWN (P2) is allowed access through interface fc1/1 (F1).
- A nWWN (N1) is allowed access through interface fc1/2 (F2).
- Any WWN is allowed access through interface fc1/3 (F3).
- A nWWN (N3) is allowed access through any interface.
- A pWWN (P3) is allowed access through interface fc1/4 (F4).
- A sWWN (S1) is allowed access through interface fc1/10-13 (F10 to F13).
- A pWWN (P10) is allowed access through interface fc1/11 (F11).

Table 41-2 summarizes the port security authorization results for this active database.

Table 41-2 Authorization Results for Scenario

Scenario	Device Connection Request	Authorization	Condition	Reason
1	P1, N2, F1	Permitted	1	No conflict.
2	P2, N2, F1	Permitted	1	No conflict.
3	P3, N2, F1	Denied	2	F1 is bound to P1/P2.
4	P1, N3, F1	Permitted	6	Wildcard match for N3.

Table 41-2 Authorization Results for Scenario (continued)

Scenario	Device Connection Request	Authorization	Condition	Reason
5	P1, N1, F3	Permitted	5	Wildcard match for F3.
6	P1, N4, F5	Denied	2	P1 is bound to F1.
7	P5, N1, F5	Denied	2	N1 is only allowed on F2.
8	P3, N3, F4	Permitted	1	No conflict.
9	S1, F10	Permitted	1	No conflict.
10	S2, F11	Denied	7	P10 is bound to F11.
11	P4, N4, F5 (auto-learn on)	Permitted	3	No conflict.
12	P4, N4, F5(auto-learn off)	Denied	4	No match.
13	S3, F5 (auto-learn on)	Permitted	3	No conflict.
14	S3, F5 (auto-learn off)	Denied	4	No match.
15	P1, N1, F6 (auto-learn on)	Denied	2	P1 is bound to F1.
16	P5, N5, F1 (auto-learn on)	Denied	7	P3 is bound to F1.
17	S3, F4 (auto-learn on)	Denied	7	P3 paired with F4.
18	S1, F3 (auto-learn on)	Permitted	5	No conflict.
19	P5, N3, F3	Permitted	6	Wildcard (*) match for F3 and N3.
20	P7, N3, F9	Permitted	6	Wildcard (*) match for N3.

Port Security Manual Configuration

To configure port security on any switch in the Cisco MDS 9000 Family, follow these steps:

Step 1 Identify the WWN of the ports that need to be secured.

Step 2 Secure the fWWN to an authorized nWWN or pWWN.

Step 3 Activate the port security database.

Step 4 Verify your configuration.

This section includes the following topics:

About WWN Identification

If you decide to manually configure port security, be sure to adhere to the following guidelines:

- Identify switch ports by the interface or by the fWWN.
- Identify devices by the pWWN or by the nWWN.
- If an Nx port is allowed to log in to SAN switch port Fx, then that Nx port can only log in through the specified Fx port.
- If an Nx port's nWWN is bound to an Fx port WWN, then all pWWNs in the Nx port are implicitly paired with the Fx port.
- TE port checking is done on each VSAN in the allowed VSAN list of the trunk port.
- All PortChannel xE ports must be configured with the same set of WWNs in the same PortChannel.
- E port security is implemented in the port VSAN of the E port. In this case the sWWN is used to secure authorization checks.
- Once activated, the config database can be modified without any effect on the active database.
- By saving the running configuration, you save the configuration database and activated entries in the active database. Learned entries in the active database are not saved.

Adding Authorized Port Pairs

After identifying the WWN pairs that need to be bound, add those pairs to the port security database.

Tip Remote switch binding can be specified at the local switch. To specify the remote interfaces, you can use either the fWWN or sWWN-interface combination.

To add authorized port pairs for port security using Fabric Manager, follow these steps:

Step 1 Expand a **VSAN** and select **Port Security** in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane shown in Figure 41-6.

Figure 41-6 Port Security Configuration

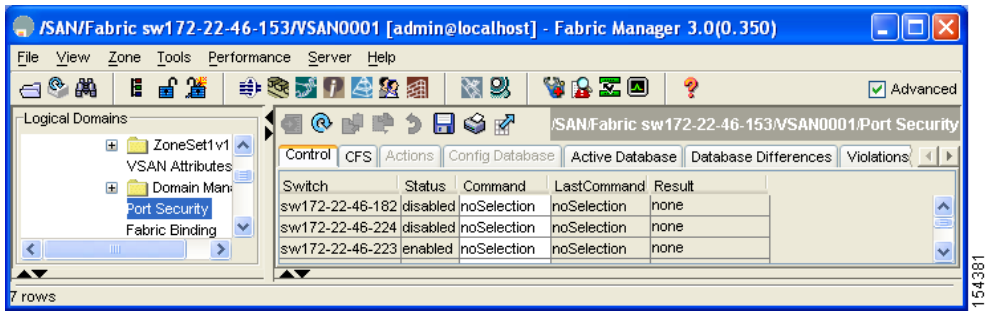

Step 2 Click the **Config Database** tab.

Step 3 Click **Create Row** to add an authorized port pair.

You see the Create Port Security dialog box in Figure 41-7.

Figure 41-7 Create Port Security Dialog Box

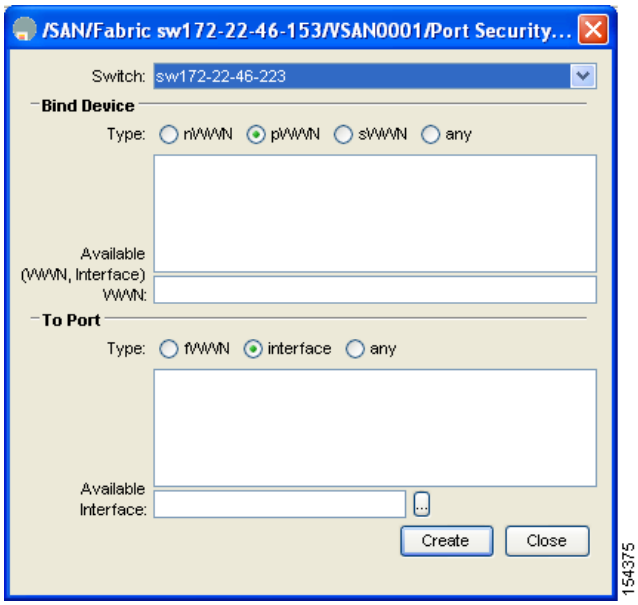

Step 4 Double-click the device from the available list for which you want to create the port security setting.

Step 5 Double-click the port from the available list to which you want to bind the device (see Figure 41-7).

Step 6 Click **Create** to create the port security setting, or click **Close** to close the Create Port Setting dialog box without adding a new port security setting (see Figure 41-7).

Step 7 Click the **CFS** tab and set the command column to **commit** on all participating switches in the VSAN.

Step 8 Click **Apply Changes** to save these changes or click **Undo Changes** to discard unsaved changes.

Deleting Port Security Setting

To delete a port security setting from the configured database on a switch, follow these steps:

Step 1 Expand a **VSAN** and select **Port Security** in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Figure 41-8 Port Security Configuration

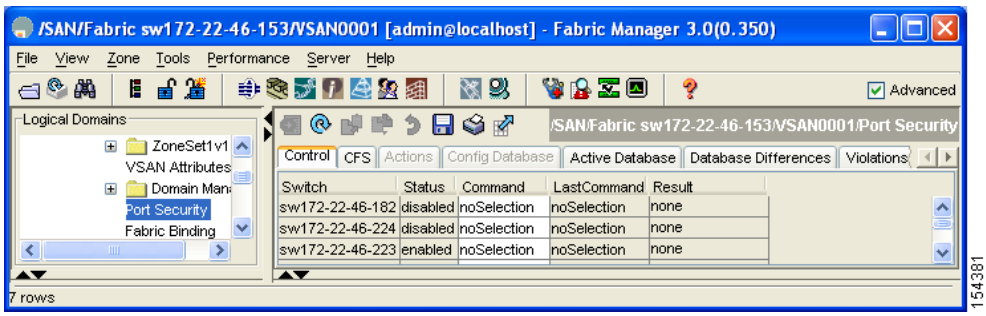

Step 2 Click the **Config Database** tab.

You see the configured port security settings for that VSAN.

Step 3 Click the row you want to delete.

Step 4 Click **Delete Row**.

You see the confirmation dialog box.

Step 5 Click **Yes** to delete the row, or click **No** to close the confirmation dialog box without deleting the row.

Step 6 Click the **CFS** tab and set the command column to **commit** on all participating switches in the VSAN.

Step 7 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

Port Security Configuration Distribution

The port security feature uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management, provide a single point of configuration for the entire fabric in the VSAN, and enforce the port security policies throughout the fabric (see Chapter 12, "Using the CFS Infrastructure").

This section contains the following topics:

Enabling Distribution

All the configurations performed in distributed mode are stored in a pending (temporary) database. If you modify the configuration, you need to commit or discard the pending database changes to the configurations. The fabric remains locked during this period. Changes to the pending database are not reflected in the configurations until you commit the changes.

Note Port Activation or deactivation and auto-learning enable or disable do not take effect until after a CFS commit if CFS distribution is enabled. Always follow any one of these operations with a CFS commit to ensure proper configuration. See the "Activation and Auto-learning Configuration Distribution" section on page 41-18.

To enable distribution using Fabric Manager, follow these steps:

Step 1 Expand a **VSAN** and select **Port Security** in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Figure 41-9 Port Security Configuration

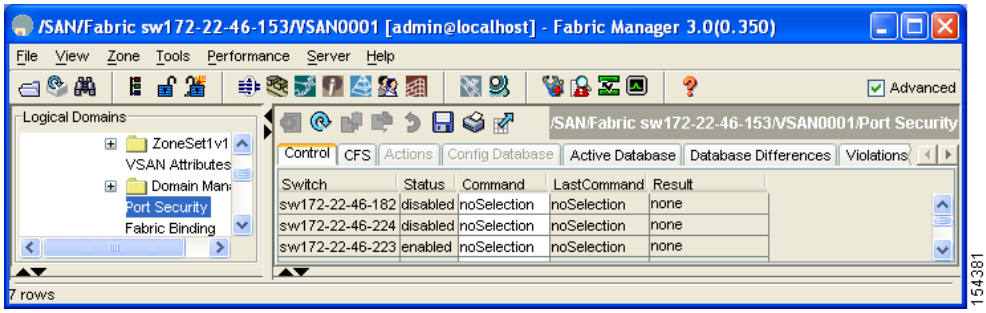

Step 2 Click the **CFS** tab.

You see the switches for that VSAN.

Step 3 Click the **Global** column and select enable or disable from the drop-down menu.

Step 4 Click **Apply Changes**.

Locking The Fabric

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

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database.

Committing the Changes

If you commit the changes made to the configurations, the configurations in the pending database are distributed to other switches. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

Activation and Auto-learning Configuration Distribution

Activation and auto-learning configurations in distributed mode are remembered merely as actions to be performed when you commit the changes in the pending database.

Learned entries are temporary and do not have any role in determining if a login is authorized or not. As such, learned entries do not participate in distribution. When you disable learning and commit the changes in the pending database, the learned entries become static entries in the active database and are distributed to all switches in the fabric. After the commit, the active database on all switches are identical and learning can be disabled.

If the pending database contains more than one activation and auto-learning configuration when you commit the changes, then the activation and auto-learning changes are consolidated and the behavior may change (see Table 41-3).

Table 41-3 Scenarios for Activation and Auto-learning Configurations in Distributed Mode

Scenario	Actions	Distribution = OFF	Distribution = ON
A and B exist in the configuration database, activation is not done and devices C,D are logged in.	1. You activate the port security database and enable auto-learning.	configuration database = {A,B} active database = {A,B, C[1], D*}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled}
	2. A new entry E is added to the configuration database.	configuration database = {A,B, E} active database = {A,B, C*, D*}	configuration database = {A,B} active database = {null} pending database = {A,B, E + activation to be enabled}
	3. You issue a commit.	Not applicable	configuration database = {A,B, E} active database = {A,B, E, C*, D*} pending database = empty
A and B exist in the configuration database, activation is not done and devices C,D are logged in.	1. You activate the port security database and enable auto-learning.	configuration database = {A,B} active database = {A,B, C*, D*}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled}
	2. You disable learning.	configuration database = {A,B} active database = {A,B, C, D}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled + learning to be disabled}
	3. You issue a commit.	Not applicable	configuration database = {A,B} active database = {A,B} and devices C and D are logged out. This is equal to an activation with auto-learning disabled. pending database = empty

1. The * (asterisk) indicates learned entries.

Tip In this case, we recommend that you perform a commit at the end of each operation: after you activate port security and after you enable auto learning.

Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the active database. See the CFS Merge Support, page 12-9 for detailed concepts.

When merging the database between two fabrics, follow these guidelines:

- Verify that the activation status and the auto-learning status is the same in both fabrics.
- Verify that the combined number of configurations for each VSAN in both databases does not exceed 2K.

Caution If you do not follow these two conditions, the merge will fail. The next distribution will forcefully synchronize the databases and the activation states in the fabric.

This section includes the following topics:

Database Interaction

Table 41-4 lists the differences and interaction between the active and configuration databases.

Table 41-4 Active and Configuration Port Security Databases

Active Database	Configuration Database
Read-only.	Read-write.
Saving the configuration only saves the activated entries. Learned entries are not saved.	Saving the configuration saves all the entries in the configuration database.
Once activated, all devices that have already logged into the VSAN are also learned and added to the active database.	Once activated, the configuration database can be modified without any effect on the active database.
You can overwrite the active database with the configured database by activating the port security database. Forcing an activation may violate the entries already configured in the active database.	You can overwrite the configuration database with the active database.

Database Scenarios

Figure 41-10 depicts various scenarios to depict the active database and the configuration database status based on port security configurations.

Figure 41-10 Port Security Database Scenarios

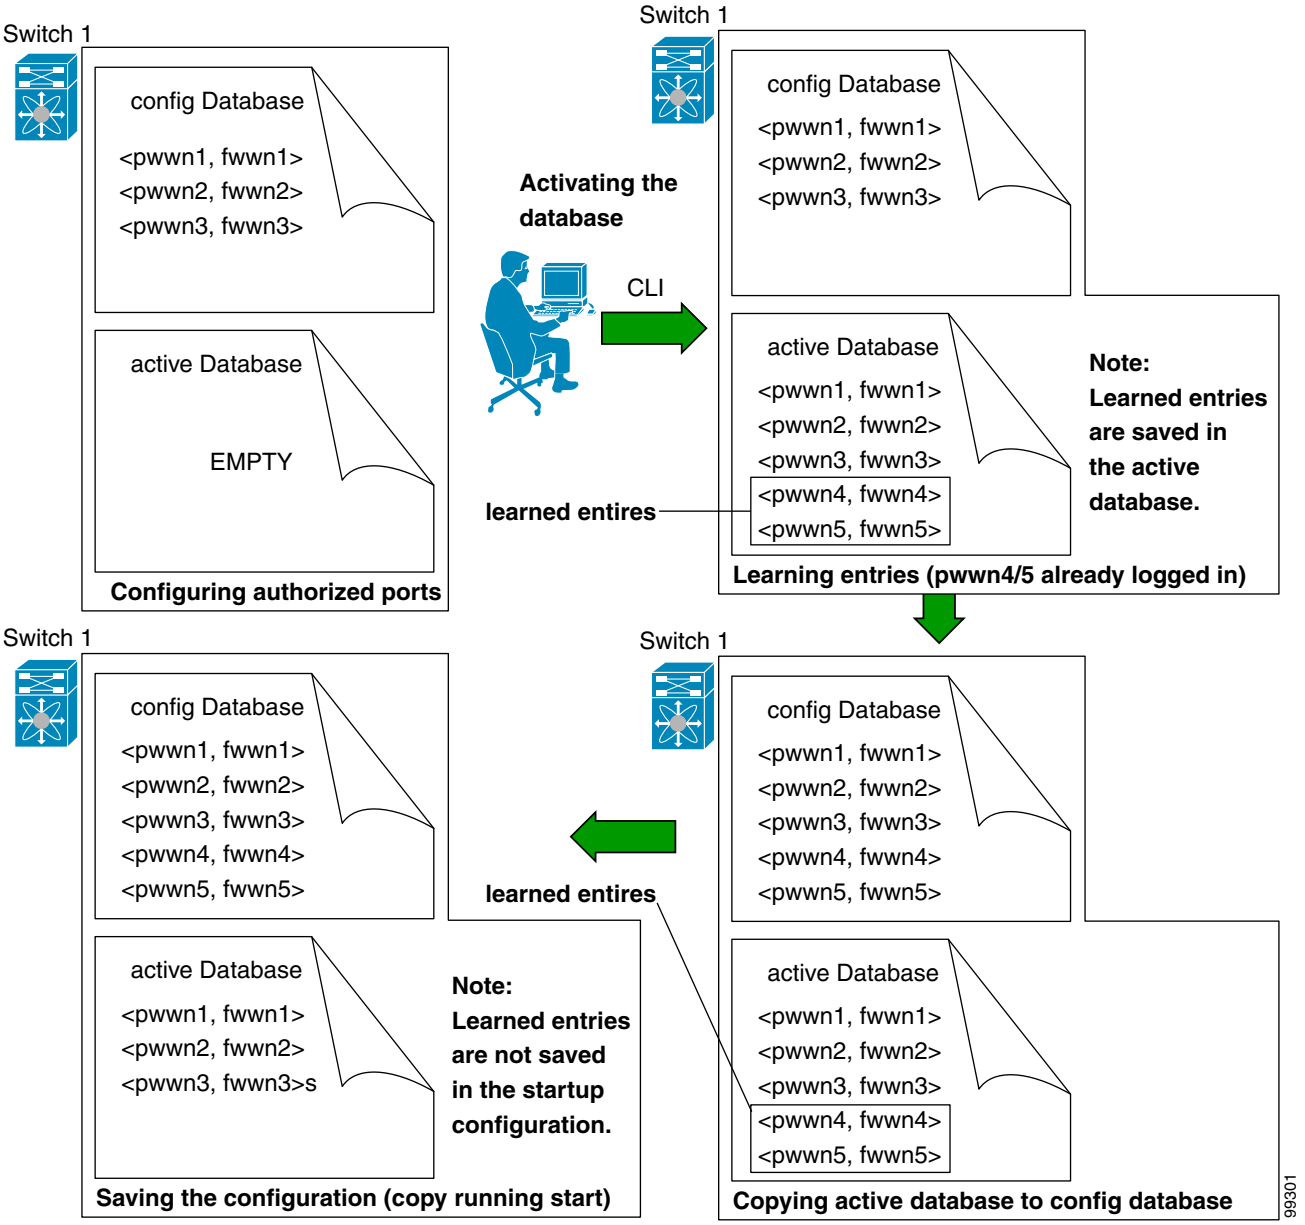

Port Security Database Copy

Tip We recommend that you copy the active database to the config database after disabling auto-learning. This action will ensure that the configuration database is in sync with the active database. If distribution is enabled, this command results in acquisition of a temporary copy (and consequently a fabric lock) of the configuration database. If you lock the fabric, you need to commit the changes to the configuration databases in all the switches.

To copy the active database to the configuration database, using Fabric Manager, follow these steps:

Step 1 Expand a **Fabric**, expand a **VSAN** and then select **Port Security** in the Logical Domains pane.

You see the Port Security information in the Information pane.

Figure 41-11 Port Security Configuration

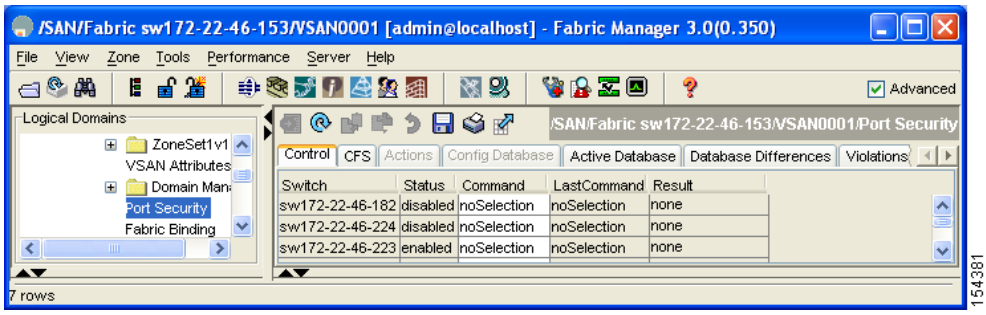

Step 2 Select the **Actions** tab. You see all the configuration databases.

Step 3 Select the appropriate configuration database and check the **Copy Active to Config** checkbox.

Step 4 Click **Apply Changes** to save your changes.

To view the differences between the active database and the configuration database using Fabric Manager, follow these steps:

Step 1 Expand a **Fabric**, expand a **VSAN** and then select **Port Security** in the Logical Domains pane.

You see the Port Security information in the Information pane.

Figure 41-12 Port Security Configuration

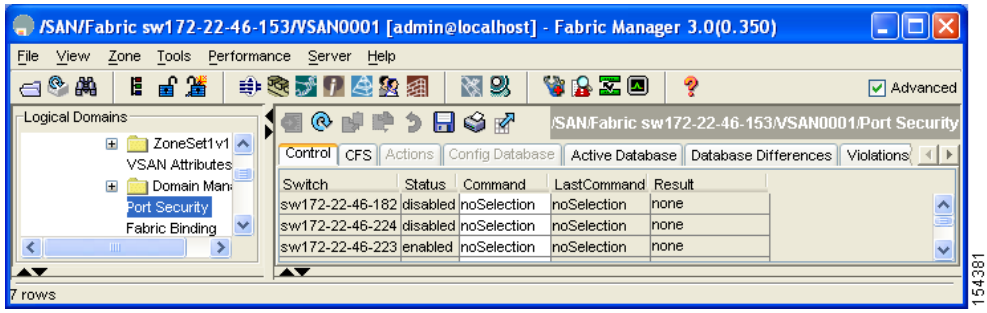

Step 2 Select the **Database Differences** tab. You see all the configuration databases.

Step 3 Select the appropriate configuration database. Select the **Active** or **Config** option to compare the differences between the selected database and the active or configuration database.

Step 4 Click **Apply Changes** to save your changes.

Port Security Database Deletion

Tip If the distribution is enabled, the deletion creates a copy of the database. An explicit deletion is required to actually delete the database.

To delete a port security database using Fabric Manager, follow these steps:

Step 1 Expand a **Fabric**, expand a **VSAN** and then select **Port Security** in the Logical Domains pane.

You see the Port Security information in the Information pane. See Figure 41-13.

Figure 41-13 Port Security Configuration

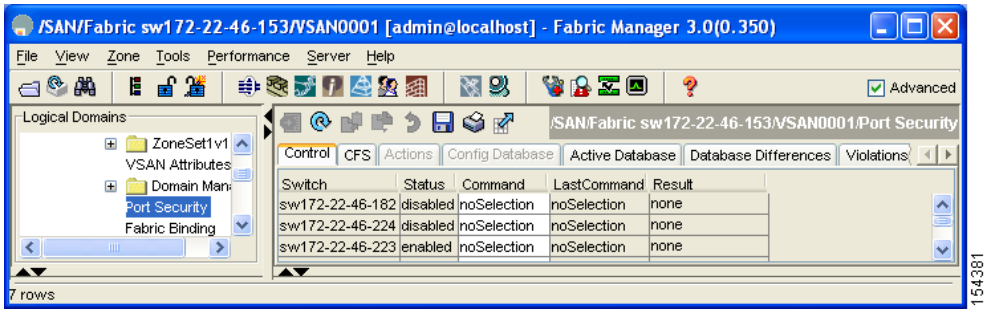

Step 2 Select the **Config Database** tab. You see all the configuration databases.

Step 3 Select the appropriate configuration database and click the **Delete Row** button.

Step 4 Click **Yes** if you want to delete the configuration database or you may click **No**.

Port Security Database Cleanup

To clear all existing statistics from the port security database for a specified VSAN using Fabric Manager, follow the steps below:

Step 1 Expand a **Fabric**, expand a **VSAN** and then select **Port Security** in the Logical Domains pane.

You see the Port Security information in the Information pane.

Figure 41-14 Port Security Configuration

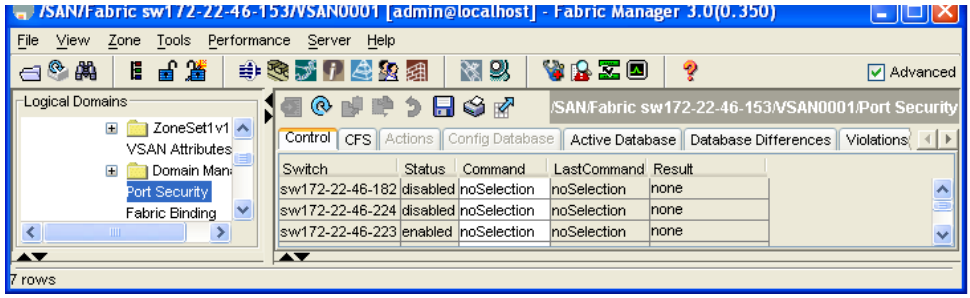

Step 2 Select the **Statistics** tab.

You see all the configuration databases.

Step 3 Select the appropriate configuration database and check the **Clear** option.

Step 4 Click **Apply Changes** to save your changes.

To clear any learned entries in the active database for a specified interface within a VSAN using Fabric Manager, follow the steps below:

Step 1 Expand a **Fabric**, expand a **VSAN** and then select **Port Security** in the Logical Domains pane.

You see the Port Security information in the Information pane.

Step 2 Select the **Actions** tab. You see all the configuration databases.

Step 3 Select the appropriate configuration database and check the **AutoLearn** option.

Step 4 Click **Apply Changes** to save your changes.

Note You can clear the Statistics and the AutoLearn option only for switches that are local and which do not acquire locks. Also, learned entries are only local to the switch and do not participate in distribution.

Default Settings

Table 41-5 lists the default settings for all port security features in any switch.

Table 41-5 Default Security Settings

Parameters	Default
Autolearn	Enabled if port security is enabled.
Port security	Disabled.
Distribution	Disabled. **Note** Enabling distribution enables it on all VSANs in the switch.

CHAPTER 42

Configuring Fabric Binding

This chapter describes the fabric binding feature provided in the Cisco MDS 9000 Family of directors and switches. It includes the following sections:

About Fabric Binding

The fabric binding feature ensures ISLs are only enabled between specified switches in the fabric binding configuration. Fabric binding is configured on a per-VSAN basis.

This feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. It uses the Exchange Fabric Membership Data (EFMD) protocol to ensure that the list of authorized switches is identical in all switches in the fabric.

This section has the following topics:

Licensing Requirements

Fabric binding requires that you install either the MAINFRAME_PKG license or the ENTERPRISE_PKG license on your switch.

See Chapter 10, “Obtaining and Installing Licenses” for more information on license feature support and installation.

Port Security Versus Fabric Binding

Port security and fabric binding are two independent features that can be configured to complement each other. Table 42-1 compares the two features.

Table 42-1 Fabric Binding and Port Security Comparison

Fabric Binding	Port Security
Uses a set of sWWNs and a persistent domain ID.	Uses pWWNs/nWWNs or fWWNs/switch WWNs.
Binds the fabric at the switch level.	Binds devices at the interface level.
Authorizes only the configured sWWN stored in the fabric binding database to participate in the fabric.	Allows a preconfigured set of Fibre Channel devices to logically connect to a SAN ports. The switch port, identified by a WWN or interface number, connects to a Fibre Channel device (a host or another switch), also identified by a WWN. By binding these two devices, you lock these two ports into a group (or list).
Requires activation on a per VSAN basis.	Requires activation on a per VSAN basis.
Allows specific user-defined switches that are allowed to connect to the fabric, regardless of the physical port to which the peer switch is connected.	Allows specific user-defined physical ports to which another device can connect.
Does not learn about switches that are logging in.	Learns about switches or devices that are logging in if learning mode is enabled.
Cannot be distributed by CFS and must be configured manually on each switch in the fabric.	Can be distributed by CFS.

Port-level checking for xEports is as follows:

- The switch login uses both port security binding and fabric binding for a given VSAN.
- Binding checks are performed on the port VSAN as follows:
  - E port security binding check on port VSAN
  - TE port security binding check on each allowed VSAN

While port security complements fabric binding, they are independent features and can be enabled or disabled separately.

Fabric Binding Enforcement

To enforce fabric binding, configure the switch world wide name (sWWN) to specify the xE port connection for each switch. Enforcement of fabric binding policies are done on every activation and when the port tries to come up. In a FICON VSAN, the fabric binding feature requires all sWWNs connected to a switch and their persistent domain IDs to be part of the fabric binding active database. In a Fibre Channel VSAN, only the sWWN is required; the domain ID is optional.

Note All switches in a Fibre Channel VSAN using fabric binding must be running Cisco MDS SAN-OS Release 3.0(1) or later.

To configure fabric binding in each switch in the fabric, follow these steps.

Step 1 Enable the fabric configuration feature.

Step 2 Configure a list of sWWNs and their corresponding domain IDs for devices that are allowed to access the fabric.

Step 3 Activate the fabric binding database.

Step 4 Copy the fabric binding active database to the fabric binding config database.

Step 5 Save the fabric binding configuration.

Step 6 Verify the fabric binding configuration.

Fabric Binding Configuration

The fabric binding feature ensures ISLs are only enabled between specified switches in the fabric binding configuration. Fabric binding is configured on a per-VSAN basis and can only be implemented in FICON VSANs. You can still perform fabric binding configuration in a non-FICON VSAN—these configurations will only come into effect after FICON is enabled.

This feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. It uses the Exchange Fabric Membership Data (EFMD) protocol in FICON networks to ensure that the list of authorized switches is identical in all switches in the fabric.

This section has the following topics:

- Deleting the Fabric Binding Database, page 42-39

About Fabric Binding Initiation

The fabric binding feature must be enabled in each switch in the fabric that participates in the fabric binding. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family. The configuration and verification commands for the fabric binding feature are only available when fabric binding is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.

To configure fabric binding in each switch in the fabric, follow these steps.

Step 1 Enable the fabric configuration feature.

Step 2 Configure a list of sWWNs and their corresponding domain IDs for devices that are allowed to access the fabric.

Step 3 Activate the fabric binding database.

Step 4 Save the fabric binding configuration.

Step 5 Verify the fabric binding configuration.

Enabling Fabric Binding

The fabric binding feature must be enabled in each switch in the fabric that participates in the fabric binding. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family. The configuration and verification commands for the fabric binding feature are only available when fabric binding is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.

To enable fabric binding on any participating switch using Fabric Manager, follow these steps:

Step 1 Expand the VSAN with the switches on which you want to enable fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding** (see Figure 42-1).

Figure 42-1 Fabric Binding Configuration

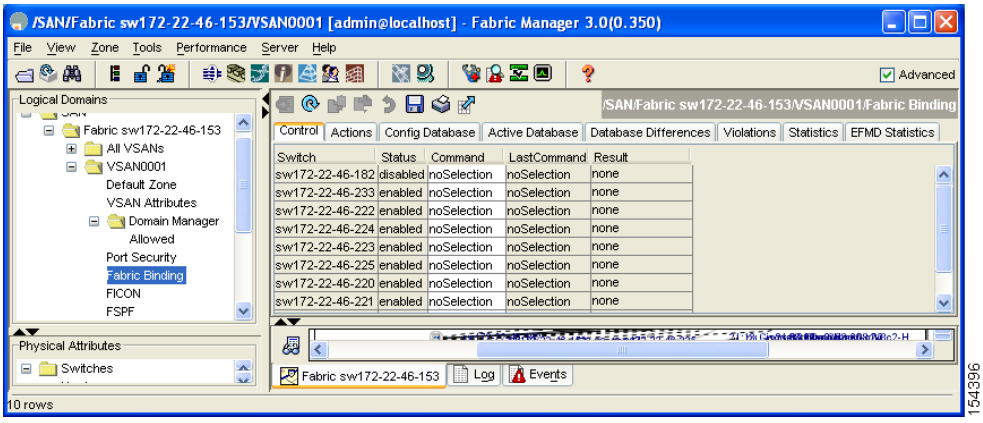

The **Control** tab is the default in the Information pane.

Step 2 Click enable or disable from the Command (see Figure 42-1) drop-down list to enable or disable Fabric Binding on the switch.

Step 3 Click **Apply Changes** to save your changes.

About Switch WWN Lists

A user-specified fabric binding list contains a list of switch WWNs (sWWNs) within a fabric. If a sWWN attempts to join the fabric, and that sWWN is not in the list or the sWWN is using a domain ID that differs from the one specified in the allowed list, the ISL between the switch and the fabric is automatically isolated in that VSAN and the switch is denied entry into the fabric.

The persistent domain ID must be specified along with the sWWN. Domain ID authorization is required in FICON VSANs where the domains are statically configured and the end devices reject a domain ID change in all switches in the fabric. Domain ID authorization is not required in Fibre Channel VSANs.

Configuring Switch WWN List

To configure a list of sWWNs and domain IDs for a FICON VSAN using Fabric Manager, follow these steps:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding** (see Figure 42-2).

Figure 42-2 Fabric Binding Configuration

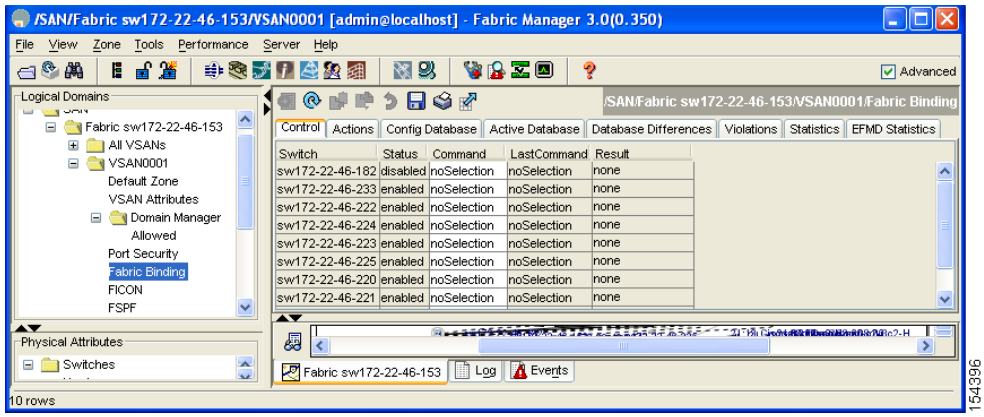

Step 2 Ensure that fabric binding is enabled for the selected VSAN. In Figure 42-2, only switch 172.22.46.182 has fabric binding enabled.

Step 3 Click the **Config Database** tab in the Information pane.

Step 4 Click **Create Row**.

You see the Create Fabric Binding dialog box in Figure 42-3.

Figure 42-3 Create Fabric Binding Dialog Box

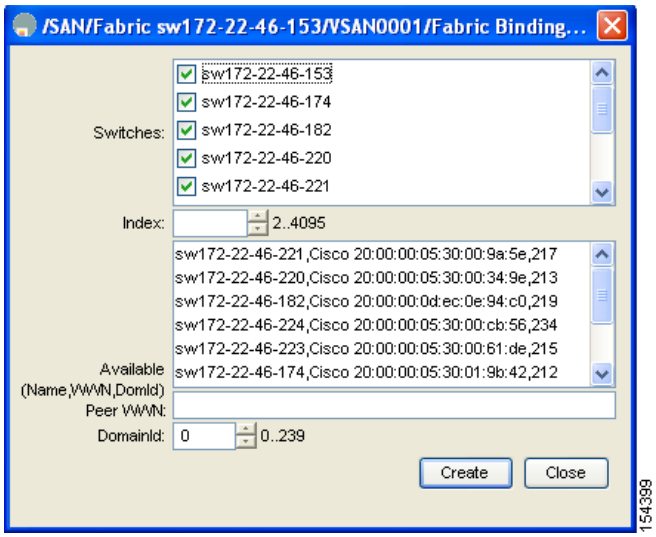

Step 5 Select the switches that you want to add.

Step 6 Add the sWWN and domain ID of a switch to the configured database list. You can add the sWWN and the domain ID of more than one switches to the configured database list.

Step 7 Click **Create**.

Fabric Binding Activation

The fabric binding feature maintains a configuration database (config-database) and an active database. The config-database is a read-write database that collects the configurations you perform. These configurations are only enforced upon activation. This activation overwrites the active database with the contents of the config- database. The active database is read-only and is the database that checks each switch that attempts to log in.

By default, the fabric binding feature is not activated. You cannot activate the fabric binding database on the switch if entries existing in the configured database conflict with the current state of the fabric. For example, one of the already logged in switches may be denied login by the config-database. You can choose to forcefully override these situations.

Note After activation, any already logged in switch that violates the current active database will be logged out, and all switches that were previously denied login because of fabric binding restrictions are reinitialized.

To copy the active database to the config database using Fabric Manager, follow these steps:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding**.

Step 2 Click the **Actions** tab in the Information pane.

Step 3 Check the **Copy Active to Config** check box.

Step 4 Click **Apply Changes** to save your changes.

Activating or Deactivating Fabric Binding

To activate, deactivate, or to force fabric binding activation using Fabric Manager, follow these steps:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding**

Step 2 Click the **Actions** tab in the Information pane.

Step 3 Click the Action drop-down menu and select **activate** or **deactivate** or **force activate** Fabric Binding on the switch.

Step 4 Click **Apply Changes** to save your changes.

The Enabled column for the switch is now **True**.

Forcing Fabric Binding Activation

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed with the activation by using the **force** option.

To forcefully activate the fabric binding database using Fabric Manager, follow these steps:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding**.

Step 2 Click the **Actions** tab in the Information pane.

Step 3 Set the Action drop-down menu to **forceActivate** for the VSAN(s) for which you want to activate fabric binding.

Step 4 Click **Apply Changes** to activate fabric binding.

The Enabled column for the switch is now **True**.

Creating a Fabric Binding Configuration

To create a fabric binding configuration using Fabric Manager, follow these steps:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding**.

Step 2 Click the **Config Database** tab in the Information pane.

You see the information in Figure 42-4.

Figure 42-4 Fabric Binding Database Configuration

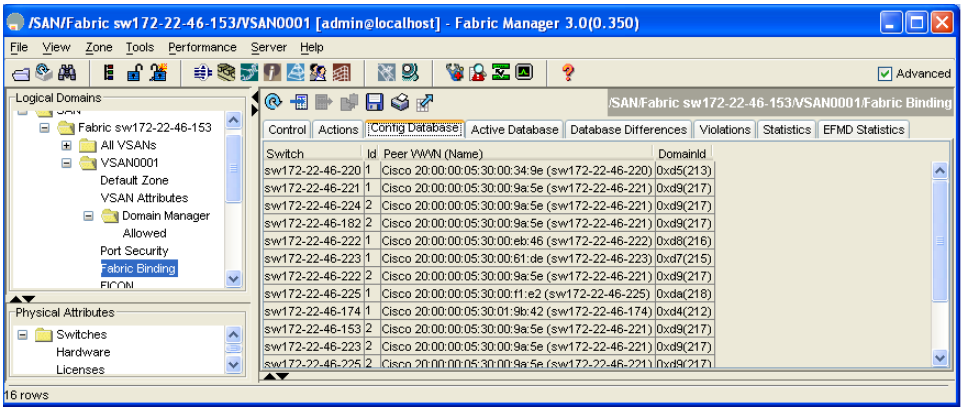

Step 3 Click **Insert Row**.

You see the Create Config Database dialog box in Figure 42-5.

Figure 42-5 Create Config Database Dialog Box

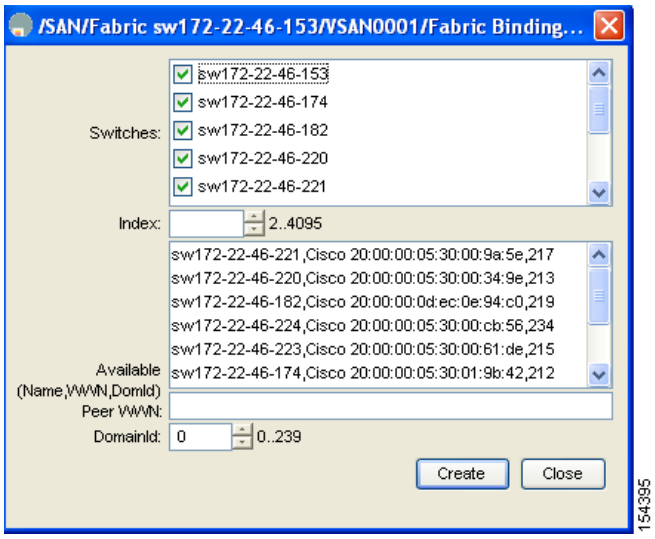

Step 4 Select switches, choose an index, and indicate the peer WWN, and the Domain ID.

Step 5 Click **Create** to create the fabric binding database configuration.

When you save the fabric binding configuration, the config database and the active database are both saved to the startup configuration and are available after a reboot.

Deleting a Fabric Binding Configuration

To delete a fabric binding configuration using Fabric Manager, follow these steps:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding**.

Step 2 Click the **Config Database** tab in the Information pane.

You see the information in Figure 42-6.

Figure 42-6 Fabric Binding Database Configuration

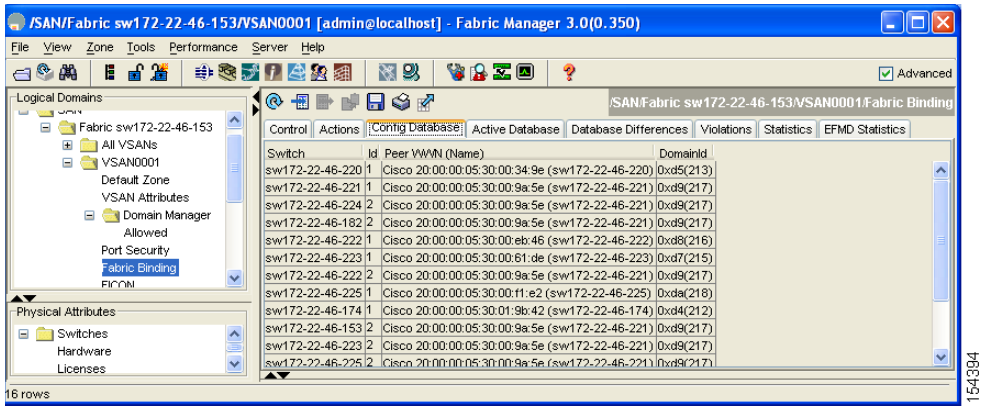

Step 3 Click in the row for the VSAN for which you want to delete the fabric binding configuration.

Step 4 Click **Delete Row** to delete the fabric binding configuration.

Copying Fabric Binding to the Configuration File

To copy the active fabric binding to the configuration file using Fabric Manager, follow these steps:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding**.

Step 2 Click the **Actions** tab in the Information pane.

You see the Fabric Binding Actions shown in Figure 42-7.

Figure 42-7 Actions Tab for Fabric Binding

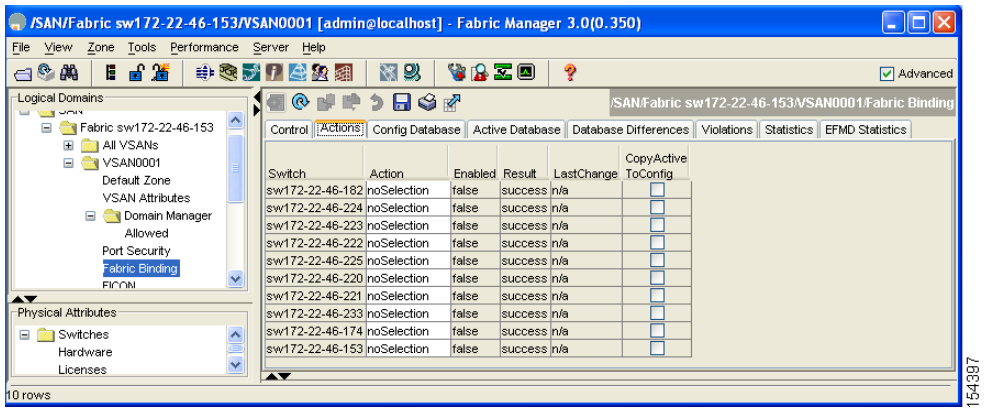

Step 3 Check the **CopyActive ToConfig** check box for the VSAN(s) for which you want to copy fabric binding (see Figure 42-7).

Step 4 Click **Apply Changes** to copy the fabric binding.

Caution You cannot deactivate or disable fabric binding in a FICON-enabled VSAN.

Clearing the Fabric Binding Statistics

To clear fabric binding statistics using Fabric Manager, follow these steps:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding**.

Step 2 Click the **Statistics** tab.

You see the statistics in the Information pane shown in Figure 42-8.

Figure 42-8 Fabric Binding Statistics

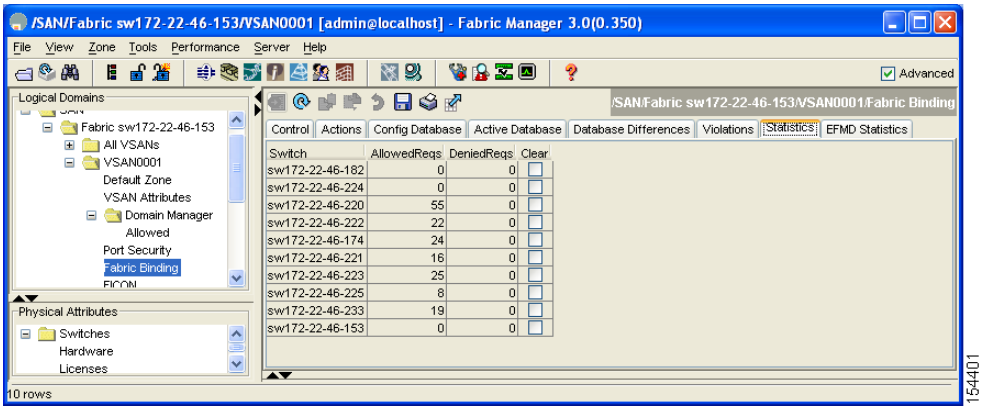

Step 3 Check the **Clear** check box for the VSAN(s) for which you want to clear statistics.

Step 4 Click **Apply Changes** to save your changes.

Viewing EFMD Statistics

To view EFMD statistics using Fabric Manager:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding**.

Step 2 Click the **EFMD Statistics** tab.

You see the EFMD statistics shown in Figure 42-9.

Figure 42-9 EFMD Statistics

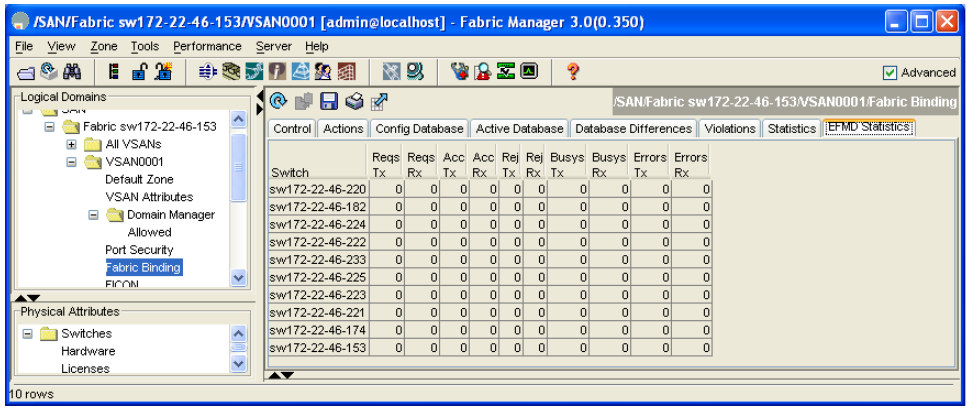

Viewing Fabric Binding Violations

To view fabric binding violations using Fabric Manager:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding**.

Step 2 Click the **Violations** tab.

You see the violations information.

Viewing Fabric Binding Active Database

To view the fabric binding active database using Fabric Manager:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding**.

Step 2 Click the **Active Database** tab.

You see the active database information shown in Figure 42-10.

Figure 42-10 Fabric Binding Active Database

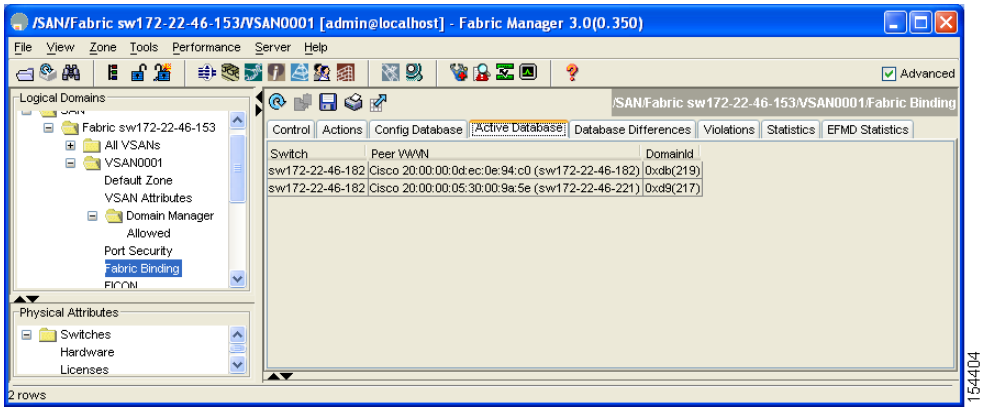

Saving Fabric Binding Configurations

When you save the fabric binding configuration, the config-database and the active database are both saved to the startup configuration and are available after a reboot.

Caution You cannot disable fabric binding in a FICON-enabled VSAN.

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding**.

Step 2 Click the **Actions** tab.

You see the Fabric Binding Actions.

Step 3 Check the **Copy Active to Config** checkbox to copy the active database to the config-database. If the configured database is empty, this action is not successful.

Step 4 Click the **Database Differences** tab to compare the database with the Config or Active database to view the differences between the active database and the config-database. Use this command to resolve conflicts.

Clearing the Fabric Binding Statistics

To clear all existing statistics from the fabric binding database for a specified VSAN using Fabric Manager, follow these steps:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding**.

Step 2 Click the **Statistics** tab in the Information pane.

You see the statistics in the Information pane

Step 3 Check the **Clear** check box.

Step 4 Click **Apply Changes** to save your changes.

Deleting the Fabric Binding Database

To delete the configured database for a specified VSAN using Fabric Manager, follow these steps:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Domain Manager** and select **Fabric Binding**.

Step 2 Click the **Config Database** tab in the Information pane.

You see the information in Figure 42-11.

Figure 42-11 Fabric Binding Database Configuration

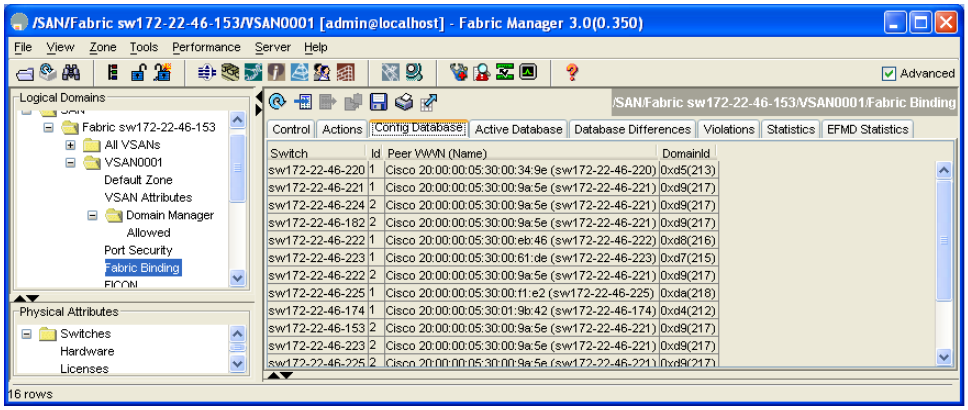

Step 3 Select the database that you want to delete.

Step 4 Click **Delete Row**.

Default Settings

Table 42-2 lists the default settings for the fabric binding feature.

Table 42-2 Default Fabric Binding Settings

Parameters	Default
Fabric binding	Disabled.

Cisco Systems

Send documentation comments to mdsfeedback-doc@cisco.com

PART 6

IP Services

CHAPTER 43

Configuring FCIP

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

Note FCIP is specific to the IPS module and is available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors.

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

Note For information on configuring Gigabit Ethernet interfaces, see the "Configuring Gigabit Ethernet Interfaces for IPv4" section on page 47-4.

This chapter includes the following sections:

About FCIP

The Fibre Channel over IP Protocol (FCIP) is a tunneling protocol that connects geographically distributed Fibre Channel storage area networks (SAN islands) transparently over IP local area networks (LANs), metropolitan area networks (MANs), and wide area networks (WANs). See Figure 43-1.

Figure 43-1 Fibre Channel SANs Connected by FCIP

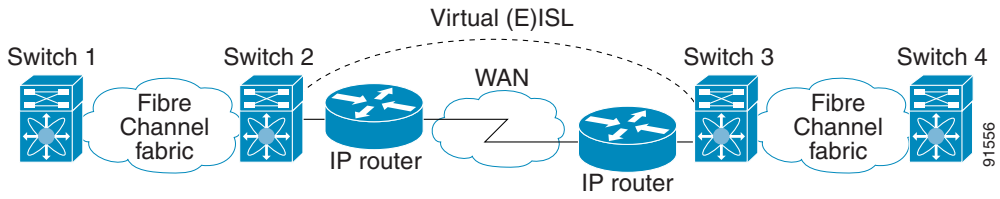

FCIP uses TCP as a network layer transport.

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

This section includes the following topics:

FCIP Concepts

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

FCIP and VE Ports

Figure 43-2 describes the internal model of FCIP with respect to Fibre Channel Inter-Switch Links (ISLs) and Cisco's extended ISLs (EISLs).

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

A virtual ISL is established over an FCIP link and transports Fibre Channel traffic. Each associated virtual ISL looks like a Fibre Channel ISL with either an E port or a TE port at each end (see Figure 43-2).

Figure 43-2 FCIP Links and Virtual ISLs

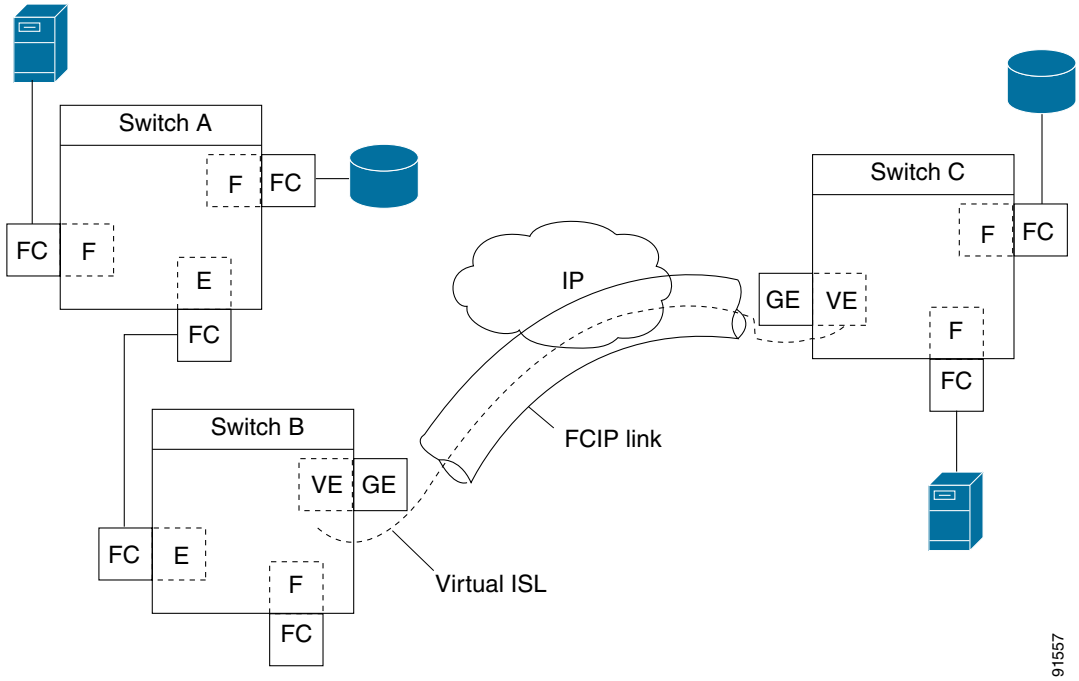

See the "E Port" section on page 18-3.

FCIP Links

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

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

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

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

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

The FCIP link is established between two peers, the VE port initialization behavior is identical to a normal E port. This behavior is independent of the link being FCIP or pure Fibre Channel, and is based on the E port discovery process (ELP, ESC).

Once the FCIP link is established, the VE port behavior is identical to E port behavior for all inter-switch communication (including domain management, zones, and VSANs). At the Fibre Channel layer, all VE and E port operations are identical.

FCIP Profiles

The FCIP profile contains information about local IP address and TCP parameters. The profile defines the following information:

- The local connection points (IP address and TCP port number).
- The behavior of the underlying TCP connections for all FCIP links that use this profile.

The FCIP profile's local IP address determines the Gigabit Ethernet port where the FCIP links terminate (see Figure 43-3).

Figure 43-3 FCIP Profile and FCIP Links

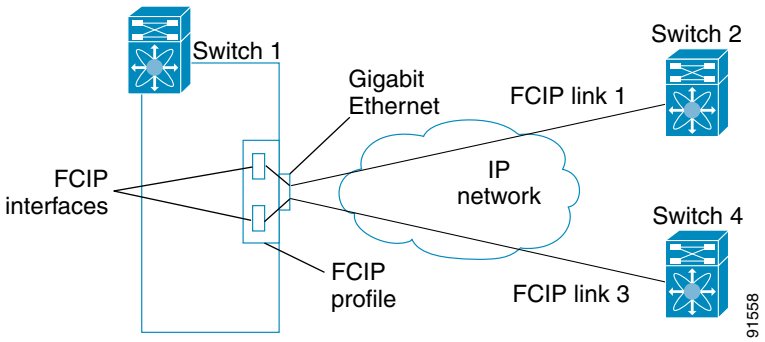

FCIP Interfaces

The FCIP interface is the local endpoint of the FCIP link and a VE port interface. All the FCIP and E port parameters are configured in context to the FCIP interface.

The FCIP parameters consist of the following:

- The FCIP profile determines which Gigabit Ethernet port initiates the FCIP links and defines the TCP connection behavior.
- Peer information.
- Number of TCP connections for the FCIP link.
- E port parameters—trunking mode and trunk allowed VSAN list.

FCIP High Availability Solutions

The following high availability solutions are available for FCIP configurations:

- Fibre Channel PortChannels, page 43-5
- FSPF, page 43-5
- VRRP, page 43-6
- Ethernet PortChannels, page 43-6

Fibre Channel PortChannels

Figure 43-4 provides an example of a PortChannel-based load balancing configuration. To perform this configuration, you need two IP addresses on each SAN island. This solution addresses link failures.

Figure 43-4 PortChannel Based Load Balancing

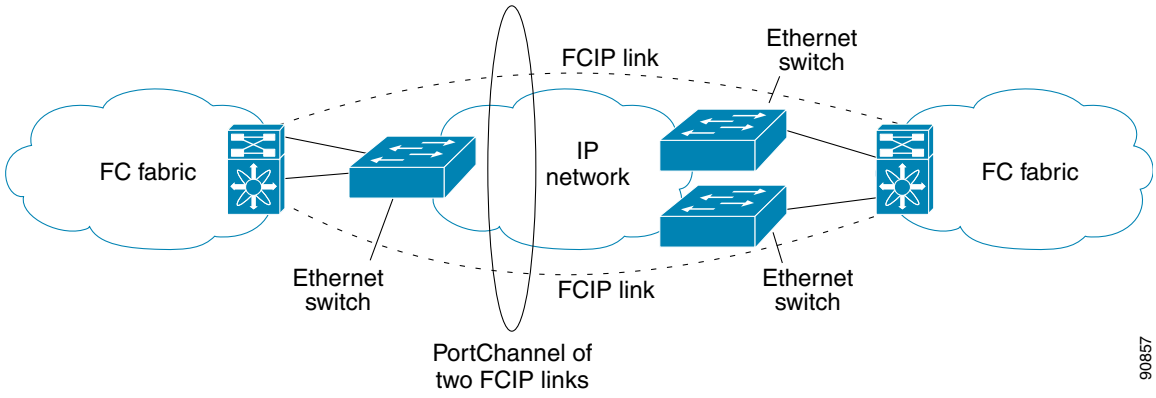

The following characteristics set Fibre Channel PortChannel solutions apart from other solutions:

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

FSPF

Figure 43-5 displays a FPSF-based load balancing configuration example. This configuration requires two IP addresses on each SAN island, and addresses IP and FCIP link failures.

Figure 43-5 FSPF-Based Load Balancing

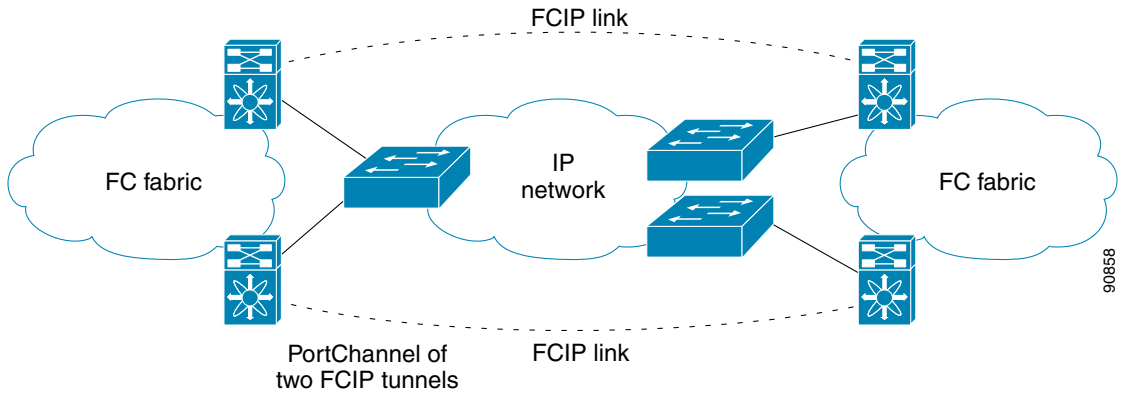

The following characteristics set FSPF solutions apart from other solutions:

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

VRRP

Figure 43-6 displays a VRRP-based high availability FCIP configuration example. This configuration requires at least two physical Gigabit Ethernet ports connected to the Ethernet switch on the island where you need to implement high availability using VRRP.

Figure 43-6 VRRP-Based High Availability

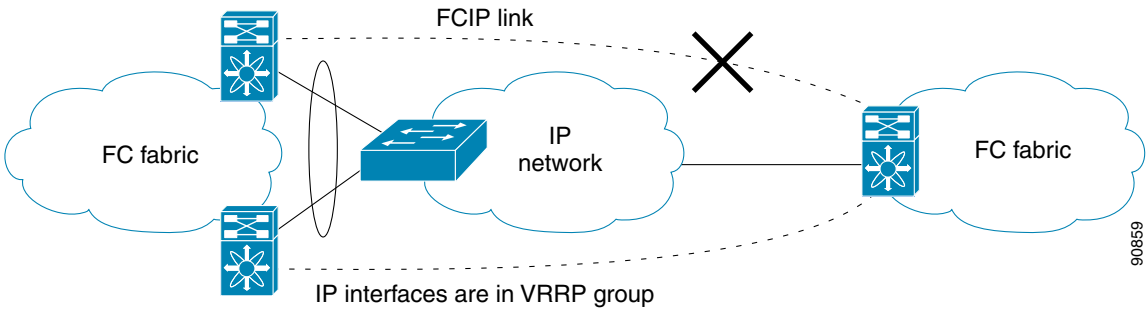

The following characteristics set VRRP solutions apart from other solutions:

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

Ethernet PortChannels

Figure 43-7 displays an Ethernet PortChannel-based high availability FCIP example. This solution addresses the problem caused by individual Gigabit Ethernet link failures.

Figure 43-7 Ethernet PortChannel-Based High Availability

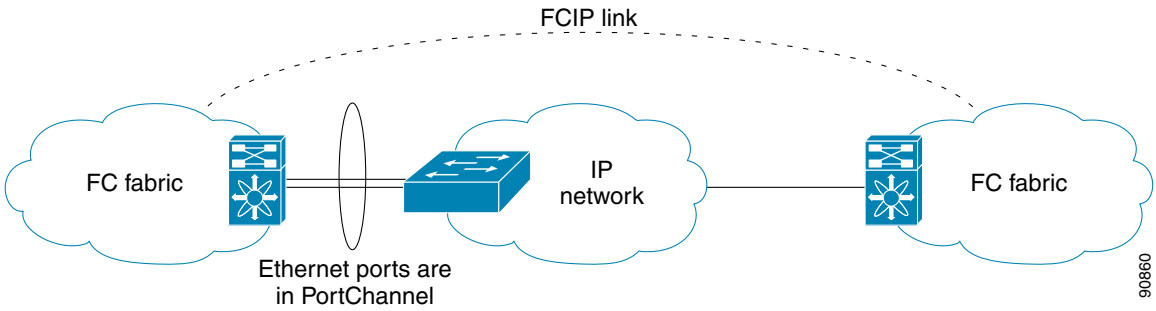

The following characteristics set Ethernet PortChannel solutions apart from other solutions:

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

Ethernet PortChannels and Fibre Channel PortChannels

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

An Ethernet PortChannel restriction only allows two contiguous IPS ports, such as ports 1–2 or 3–4, to be combined in one Ethernet PortChannel (see the "Configuring Gigabit Ethernet High Availability" section on page 47-8). This restriction only applies to Ethernet PortChannels. The Fibre Channel PortChannel (to which FCIP link can be a part of), does not have a restriction on which (E)ISL links can be combined in a Fibre Channel PortChannel as long as it passes the compatibility check (see the "Compatibility Check" section on page 21-15). The maximum number of Fibre Channel ports that can be put into a Fibre Channel PortChannel is 16 (see Figure 43-8).

Figure 43-8 PortChannels at the Fibre Channel and Ethernet Levels

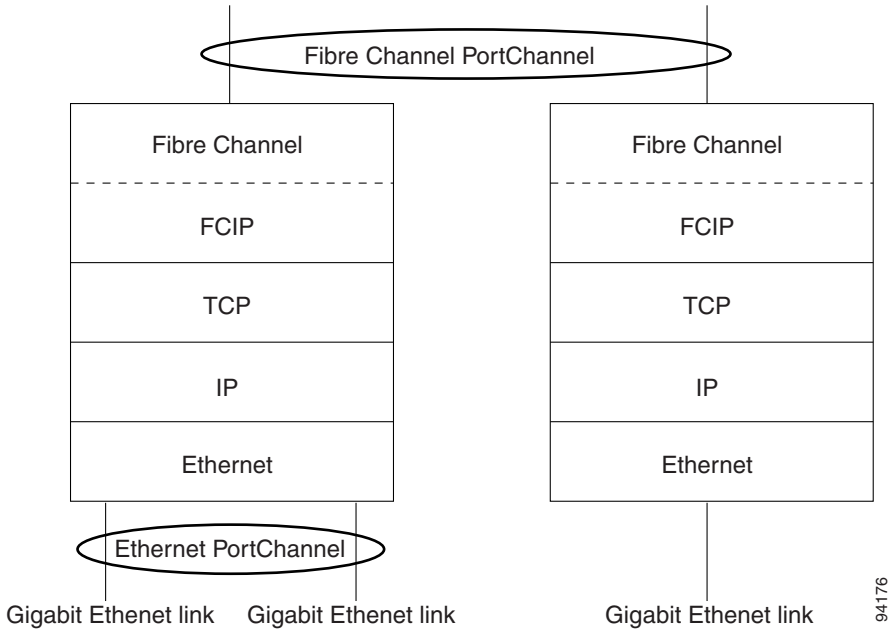

To configure Fibre Channel PortChannels, see Chapter 21, "Configuring PortChannels." To configure Ethernet PortChannels, see the "Configuring Gigabit Ethernet High Availability" section on page 47-8.

Configuring FCIP

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

Enabling FCIP

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

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

To use the FCIP feature, you need to obtain the SAN extension over IP package license (SAN_EXTN_OVER_IP or SAN_EXTN_OVER_IP_IPS4) (see Chapter 10, "Obtaining and Installing Licenses").

Using the FCIP Wizard

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

To create and manage FCIP links with Fabric Manager, use the FCIP Wizard. First verify that the IP services module is inserted in the required Cisco MDS 9000 Family switches and that the Gigabit Ethernet interfaces on these switches are connected and the connectivity verified. The steps in creating FCIP links using the FCIP Wizard are:

- Select the endpoints.
- Choose the interfaces' IP addresses.
- Specify link attributes.
- Optionally enable FCIP write acceleration or FCIP compression.

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

Step 1 Open the FCIP Wizard by clicking its icon in the Fabric Manager toolbar. Figure 43-9 shows the FCIP Wizard icon.

Figure 43-9 FCIP Wizard

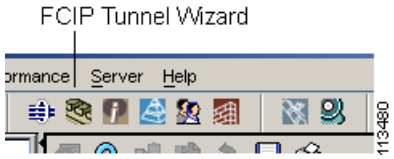

Step 2 Choose the switches that act as endpoints for the FCIP link and click **Next**.

Step 3 Choose the Gigabit Ethernet ports on each switch that will form the FCIP link.

Step 4 If both Gigabit Ethernet ports are part of MPS-14/2 modules, you can check the **Enforce IPSEC Security** check box and set the **IKE Auth Key**, as shown in Figure 43-10. See the "IPsec and IKE Terminology" section on page 39-5 for information on IPsec and IKE.

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

Figure 43-10 Enabling IPsec on an FCIP link

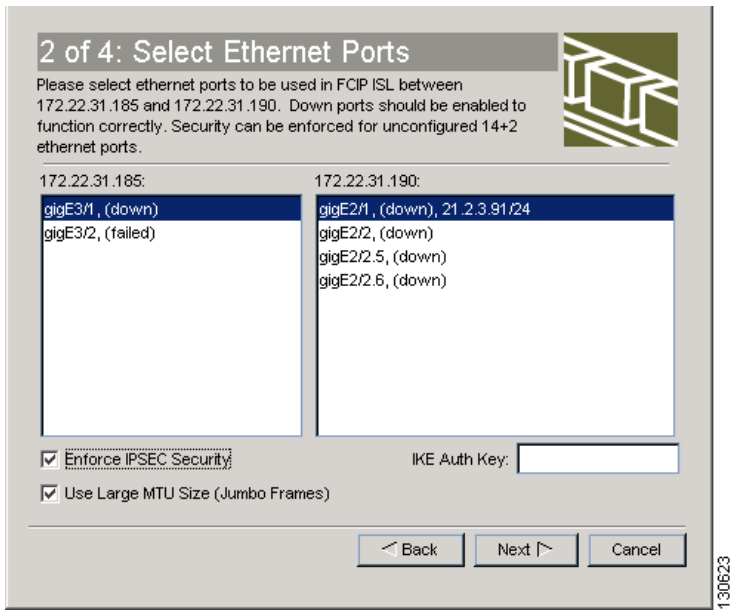

Step 5 Click **Next**. You see the TCP connection characteristics.

Step 6 Set the minimum and maximum bandwidth settings and round-trip time for the TCP c=]

Step 7 onnections on this FCIP link, as shown in Figure 43-11. You can measure the round-trip time between the Gigabit Ethernet endpoints by clicking the **Measure** button.

Figure 43-11 Specifying Tunnel Properties

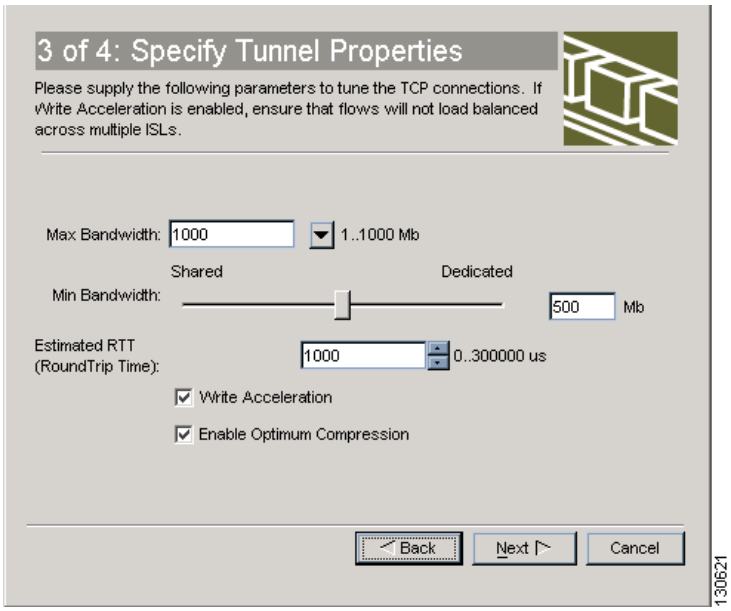

Step 8 Check the **Enable Write Acceleration** check box to enable FCIP write acceleration on this FCIP link. See the "FCIP Write Acceleration" section on page 43-23.

Step 9 Check the **Enable Optimum Compression** check box to enable IP compression on this FCIP link. See the "FCIP Compression" section on page 43-30.

Step 10 Click **Next** to configure the FCIP tunnel parameters.

Step 11 Set the **FICON Port Address** if FICON is required on this FCIP link. Click the **...** button to show the first available FICON port.

Step 12 Set the **Port VSAN** and click the **Trunk Mode** radio button for this FCIP link, as shown in Figure 43-12. See the "Checking Trunk Status" section on page 43-13.

Figure 43-12 Create FCIP ISL

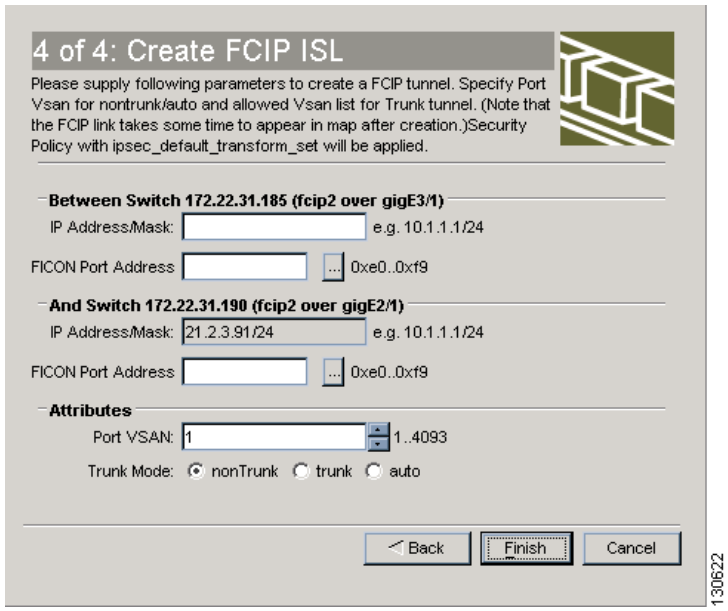

Step 13 Click **Finish** to create this FCIP link or click **Cancel** to exit the FCIP Wizard without creating an FCIP link.

Basic FCIP Configuration

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

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

Step 1 Configure the Gigabit Ethernet interface.

Step 2 Create an FCIP profile, and then assign the Gigabit Ethernet interface's IP address to the profile.

Step 3 Create an FCIP interface, and then assign the profile to the interface.

Step 4 Configure the peer IP address for the FCIP interface.

Step 5 Enable the interface.

Creating FCIP Profiles

You must assign a local IP address of a Gigabit Ethernet interface or subinterface to the FCIP profile to create an FCIP profile (see Figure 43-13).

Figure 43-13 Assigning Profiles to Each Gigabit Ethernet Interface

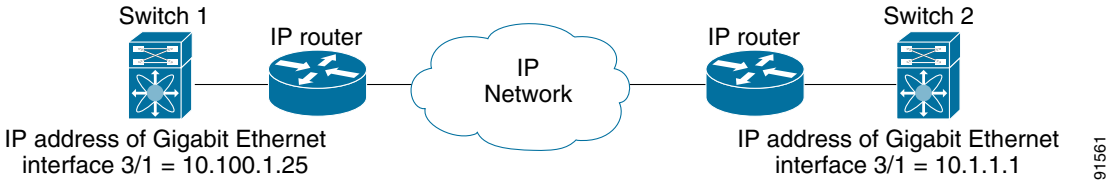

Step 1 To create an FCIP profile in switch 1, follow these steps. Verify that you are connected to a switch that contains an IPS module.

Step 2 From Fabric Manager, choose **Switches > ISLs > FCIP** in the Physical Attributes pane. From Device Manager, choose **FCIP** from the IP menu.

Step 3 Click the **Create Row** button in Fabric Manager or the **Create** button on Device Manager to add a new profile.

Step 4 Enter the profile ID in the ProfileId field.

Step 5 Enter the IP address of the interface to which you want to bind the profile.

Step 6 Modify the optional TCP parameters, if desired. Refer to Fabric Manager Online Help for explanations of these fields

Step 7 Optionally, click the **Tunnels** tab and modify the remote IP address in the Remote IPAddress field for the endpoint to which you want to link.

Step 8 Enter the optional parameters, if desired. See the "Advanced FCIP Interface Configuration" section on page 43-17.

Step 9 Click **Apply Changes** icon to save these changes or click the **Undo Changes** icon to discard any unsaved changes.

Creating FCIP Links

When two FCIP link endpoints are created, an FCIP link is established between the two IPS modules or MPS-14/2 modules. To create an FCIP link, assign a profile to the FCIP interface and configure the peer information. The peer IP switch information initiates (creates) an FCIP link to that peer switch (see Figure 43-14).

Figure 43-14 Assigning Profiles to Each Gigabit Ethernet Interface

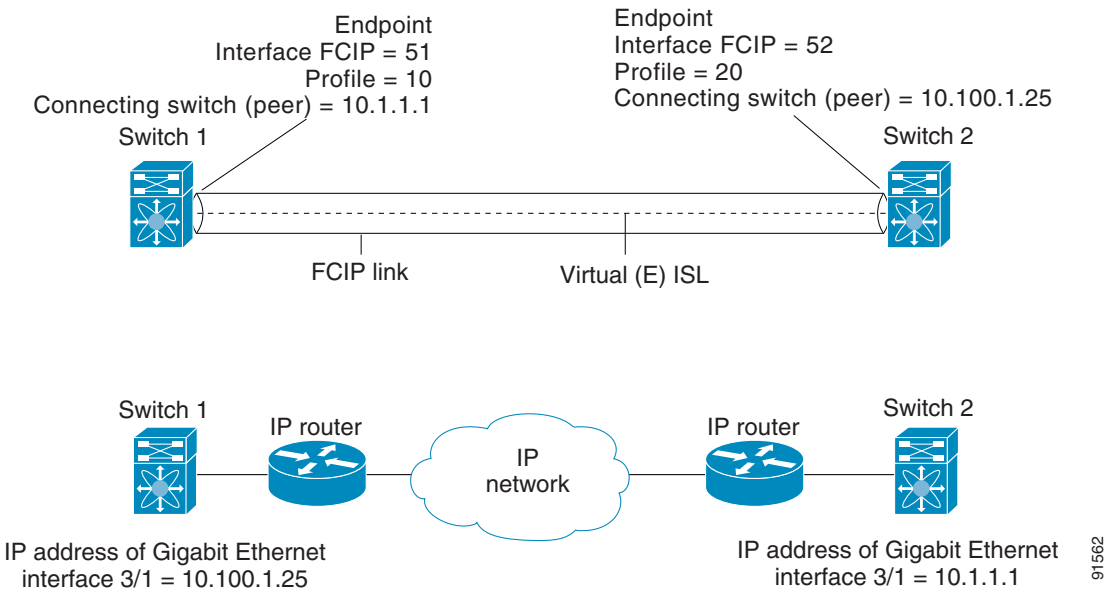

Verifying Interfaces and Extended Link Protocol

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

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Select **FCIP** from the Interface menu.

Step 3 Click the **Interfaces** tab if it is not already selected. You see the FCIP Interfaces dialog box.

Step 4 Click the **ELP** tab if it is not already selected. You see the FCIP ELP dialog box.

Checking Trunk Status

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

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Select **FCIP** from the IP menu.

Step 3 Click the **Trunk Config** tab if it is not already selected. You see the FCIP Trunk Config dialog box. This shows the status of the interface.

Step 4 Click the **Trunk Failures** tab if it is not already selected. You see the FCIP Trunk Failures dialog box.

Advanced FCIP Profile Configuration

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

- Configuring TCP Listener Ports, page 43-14
- Configuring TCP Parameters, page 43-14

Configuring TCP Listener Ports

The default TCP port for FCIP is 3225.

Configuring TCP Parameters

You can control TCP behavior in a switch by configuring the following TCP parameters.

- Minimum Retransmit Timeout, page 43-14
- Keepalive Timeout, page 43-14
- Maximum Retransmissions, page 43-15
- Path MTUs, page 43-15
- Selective Acknowledgments, page 43-15
- Window Management, page 43-15
- Monitoring Congestion, page 43-16
- Estimating Maximum Jitter, page 43-16
- Buffer Size, page 43-16

Minimum Retransmit Timeout

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

Keepalive Timeout

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

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

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

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

Maximum Retransmissions

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

Path MTUs

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

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

Selective Acknowledgments

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

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

Window Management

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

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

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

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

Monitoring Congestion

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

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

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

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

Note The default burst size is 50 KB.

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

Estimating Maximum Jitter

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

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

The default value is 1000 microseconds for FCIP interfaces.

Buffer Size

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

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

Advanced FCIP Interface Configuration

You can establish connection to a peer by configuring one or more of the following options for the FCIP interface:

- Configuring Peers, page 43-17
- Active Connections, page 43-19
- Number of TCP Connections, page 43-19
- Time Stamp Control, page 43-19
- FCIP B Port Interoperability Mode, page 43-20
- Quality of Service, page 43-22

To establish a peer connection, you must first create the FCIP interface.

Configuring Peers

To establish an FCIP link with the peer, you can use one of two options:

- Peer IP address—Configures both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.
- Special frames—Configures one end of the FCIP link when security gateways are present in the IP network. Optionally, you can also use the switch WWN (sWWN) and profile ID along with the IP address.

Peer IP Address

The basic FCIP configuration uses the peer's IP address to configure the peer information. You can also specify the peer's port number to configure the peer information. If you do not specify a port, the default 3225 port number is used to establish connection. You can specify an IPv4 address or an IPv6 address.

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

Step 1 From Fabric Manager, expand **ISLs** and select **FCIP** in the Physical Attributes pane.

You see the FCIP profiles and links in the Information pane.

From Device manager, choose **IP > FCIP**.

You see the FCIP dialog box.

Step 2 Click the **Tunnels** tab. You see the FCIP link information.

Step 3 Click the **Create Row** icon in Fabric Manager or the **Create** button in Device Manager. You see the FCIP Tunnels dialog box.

Step 4 Set the ProfileID and TunnelID fields.

Step 5 Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.

Step 6 Check the **PassiveMode** check box if you do not want this end of the link to initiate a TCP connection.

Step 7 Optionally set the NumTCPCon field to the number of TCP connections from this FCIP link.

Step 8 Optionally, check the **Enable** check box in the Time Stamp section and set the Tolerance field.

Step 9 Optionally set the other fields in this dialog box and click **Create** to create this FCIP link.

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

Step 1 From Fabric Manager, choose **ISLs > FCIP** from the Physical Attributes pane. You see the FCIP profiles and links in the Information pane.

From Device manager, choose **IP > FCIP**. You see the FCIP dialog box.

Step 2 Click the **Tunnels** tab. You see the FCIP link information.

Step 3 Click the **Create Row** icon in Fabric Manager or the **Create** button in Device Manager. You see the FCIP Tunnels dialog box.

Step 4 Set the ProfileID and TunnelID fields.

Step 5 Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.

Step 6 Check the **PassiveMode** check box if you do not want this end of the link to initiate a TCP connection.

Step 7 Optionally set the NumTCPCon field to the number of TCP connections from this FCIP link.

Step 8 Optionally, check the **Enable** check box in the Time Stamp section and set the Tolerance field.

Step 9 Optionally set the other fields in this dialog box and click **Create** to create this FCIP link.

Special Frames

You can alternatively establish an FCIP link with a peer using an optional protocol called *special frames*. When special frames are enabled, the peer IP address (and optionally the port or the profile ID) only needs to be configured on one end of the link. Once the connection is established, a special frame is exchanged to discover and authenticate the link.

By default, the special frame feature is disabled. You must enable special frames on the interfaces on both peers to establish the FCIP link.

Note Refer to the Fibre Channel IP standards for further information on special frames.

Tip Special frame negotiation provides an additional authentication security mechanism because the link validates the WWN of the peer switch.

To enable special frames, follow these steps

Step 1 From Fabric Manager, choose **ISLs > FCIP** from the Physical Attributes pane. You see the FCIP profiles and links in the Information pane.

From Device manager, choose **IP > FCIP**. You see the FCIP dialog box.

Step 2 Click the **Tunnels** tab. You see the FCIP link information.

Step 3 Click the **Create Row** icon in Fabric Manager or the **Create** button in Device Manager. You see the FCIP Tunnels dialog box.

Step 4 Set the ProfileID and TunnelID fields.

Step 5 Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.

Step 6 Check the **PassiveMode** check box if you do not want this end of the link to initiate a TCP connection.

Step 7 Optionally set the NumTCPCon field to the number of TCP connections from this FCIP link.

Step 8 Check the **Enable** check box in the Special Frames section of the dialog box and set the RemoteWWN and the RemoteProfileID fields.

Step 9 Optionally set the other fields in this dialog box and click **Create** to create this FCIP link.

Active Connections

You can configure the required mode for initiating a TCP connection. By default, active mode is enabled to actively attempt an IP connection. If you enable the passive mode, the switch does not initiate a TCP connection and merely waits for the peer to connect to it.

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

Number of TCP Connections

You can specify the number of TCP connections from an FCIP link. By default, the switch tries two (2) TCP connections for each FCIP link. You can configure one or two TCP connections. For example, the Cisco PA-FC-1G Fibre Channel port adapter, which has only one (1) TCP connection, interoperates with any switch in the Cisco MDS 9000 Family. One TCP connection is within the specified limit. If the peer initiates one TCP connection, and your MDS switch is configured for two TCP connections, then the software handles it gracefully and moves on with just one connection.

Time Stamp Control

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

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

Note The default value for packet acceptance is 2000 microseconds.
If the **time-stamp** option is enabled, be sure to configure NTP on both switches (see the "NTP Configuration" section on page 11-4).

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

FCIP B Port Interoperability Mode

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as Cisco's PA-FC-1G Fibre Channel port adapter and the SN 5428-2 storage router, implement a bridge port model to connect geographically dispersed fabrics. This model uses B port as described in the T11 Standard FC-BB-2. Figure 43-15 depicts a typical SAN extension over an IP network.

Figure 43-15 FCIP B Port and Fibre Channel E Port

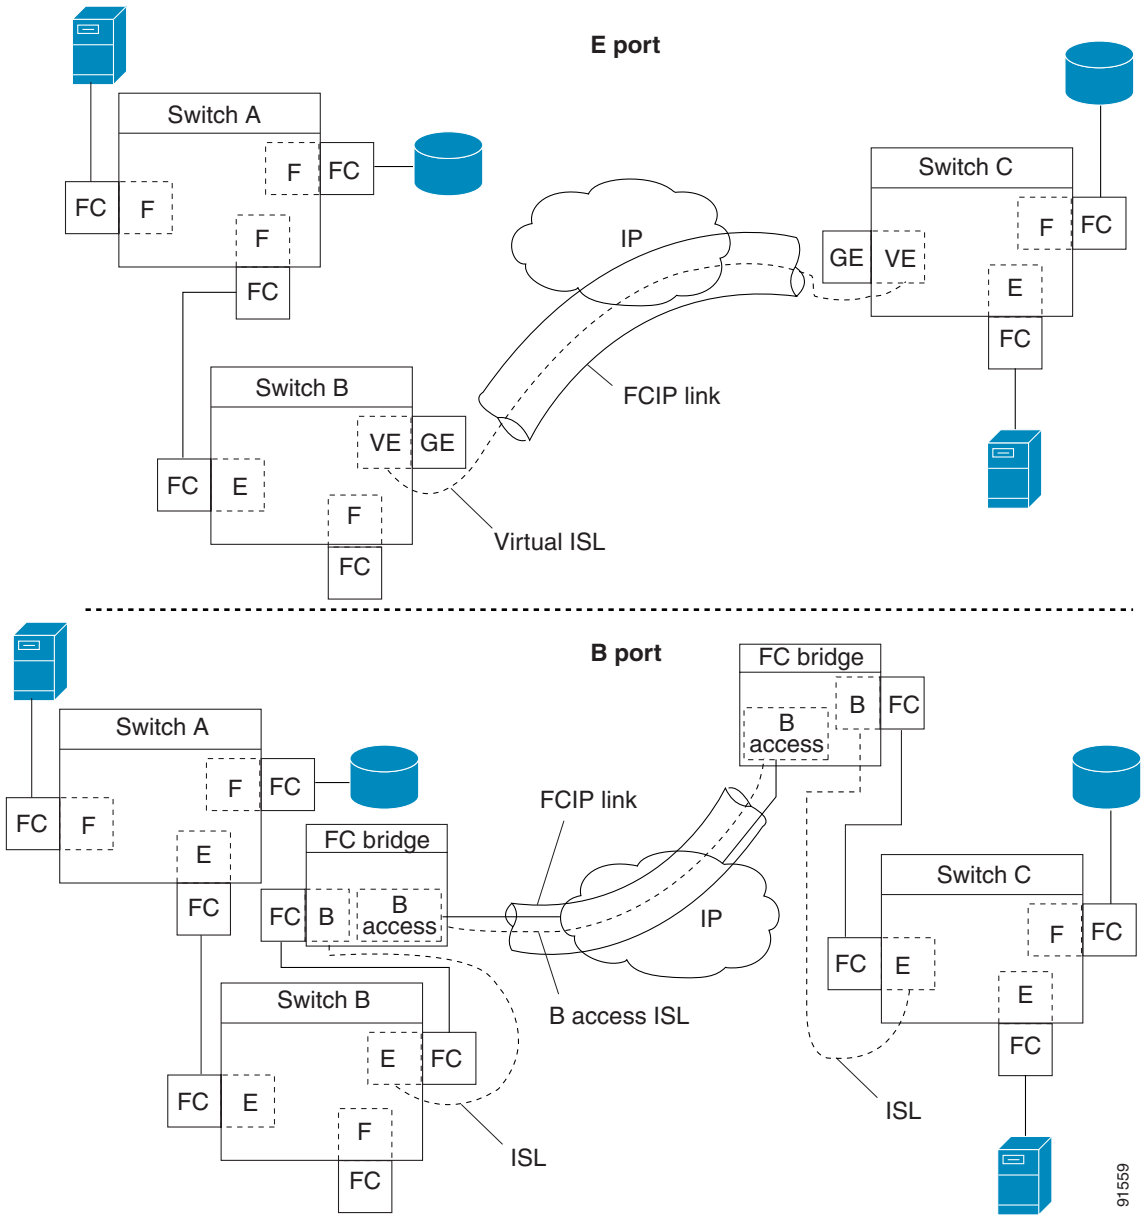

B ports bridge Fibre Channel traffic from a local E port to a remote E port without participating in fabric-related activities such as principal switch election, domain ID assignment, and Fibre Channel fabric shortest path first (FSPF) routing. For example, Class F traffic entering a SAN extender does not

interact with the B port. The traffic is transparently propagated (bridged) over a WAN interface before exiting the remote B port. This bridge results in both E ports exchanging Class F information that ultimately leads to normal ISL behavior such as fabric merging and routing.

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

The IPS module and MPS-14/2 module support FCIP links that originate from a B port SAN extender device by implementing the B access ISL protocol on a Gigabit Ethernet interface. Internally, the corresponding virtual B port connects to a virtual E port that completes the end-to-end E port connectivity requirement (see Figure 43-16).

Figure 43-16 FCIP Link Terminating in a B Port Mode

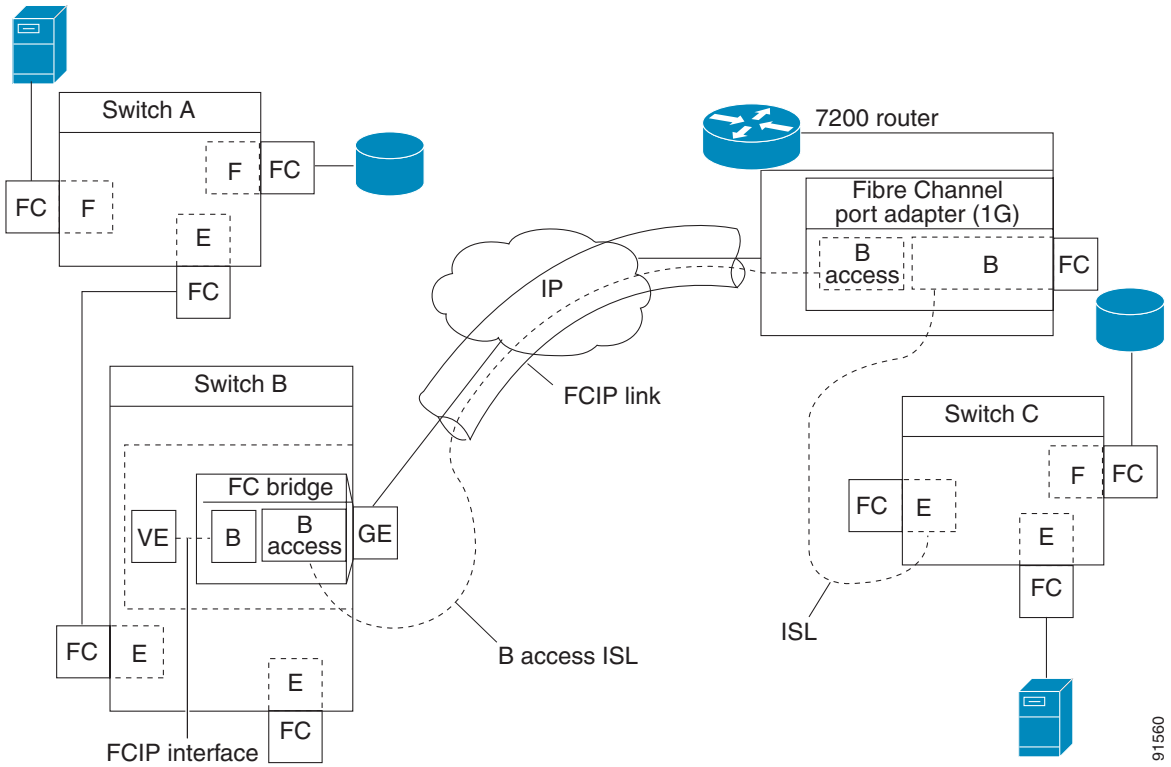

The B port feature in the IPS module and MPS-14/2 module allows remote B port SAN extenders to communicate directly with a Cisco MDS 9000 Family switch, therefore eliminating the need for local bridge devices.

Configuring B Ports

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

To enable B port mode, follow these steps:

Step 1 From Fabric Manager, choose **ISLs > FCIP** from the Physical Attributes pane. You see the FCIP profiles and links in the Information pane.

From Device manager, choose **IP > FCIP**. You see the FCIP dialog box.

Step 2 Click the **Tunnels** tab. You see the FCIP link information.

Step 3 Click the **Create Row** icon in Fabric Manager or the **Create** button in Device Manager. You see the FCIP Tunnels dialog box.

Step 4 Set the ProfileID and TunnelID fields.

Step 5 Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.

Step 6 Check the **PassiveMode** check box if you do not want this end of the link to initiate a TCP connection.

Step 7 Optionally set the NumTCPCon field to the number of TCP connections from this FCIP link.

Step 8 Check the **Enable** check box in the B Port section of the dialog box and optionally check the **KeepAlive** check box if you want a response sent to an ELS Echo frame received from the FCIP peer.

Step 9 Optionally set the other fields in this dialog box and click **Create** to create this FCIP link.

Quality of Service

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

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

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

Configuring E Ports

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

- An FCIP interface can be a member of any VSAN (see Chapter 23, “Configuring and Managing VSANs”).
- Trunk mode and trunk allowed VSANs (see Chapter 20, “Configuring Trunking”).
- PortChannels (see Chapter 21, “Configuring PortChannels”):
  - Multiple FCIP links can be bundled into a Fibre Channel PortChannel.
  - FCIP links and Fibre Channel links cannot be combined in one PortChannel.
- FSPF (see Chapter 28, “Configuring Fibre Channel Routing Services and Protocols”).

- Fibre Channel domains (fcdomains) (see Chapter 22, "Configuring Domain Parameters.").
- Importing and exporting the zone database from the adjacent switch (see Chapter 26, "Configuring and Managing Zones").

Advanced FCIP Features

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

- FCIP Write Acceleration, page 43-23
- FCIP Tape Acceleration, page 43-25
- FCIP Compression, page 43-30

FCIP Write Acceleration

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

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

In Figure 43-17, the WRITE operation performed without write acceleration requires two round trip transfers (RTT), while the WRITE operation with write acceleration only requires one RTT. The maximum sized Transfer Ready is sent from the host side of the FCIP link back to the host before the WRITE operation reaches the target. This enables the host to start sending the write data without waiting for the long latency over the FCIP link of the WRITE operation and Transfer Ready. It also eliminates the delay caused by multiple Transfer Readys needed for the exchange going over the FCIP link.

Figure 43-17 FCIP Link Write Acceleration

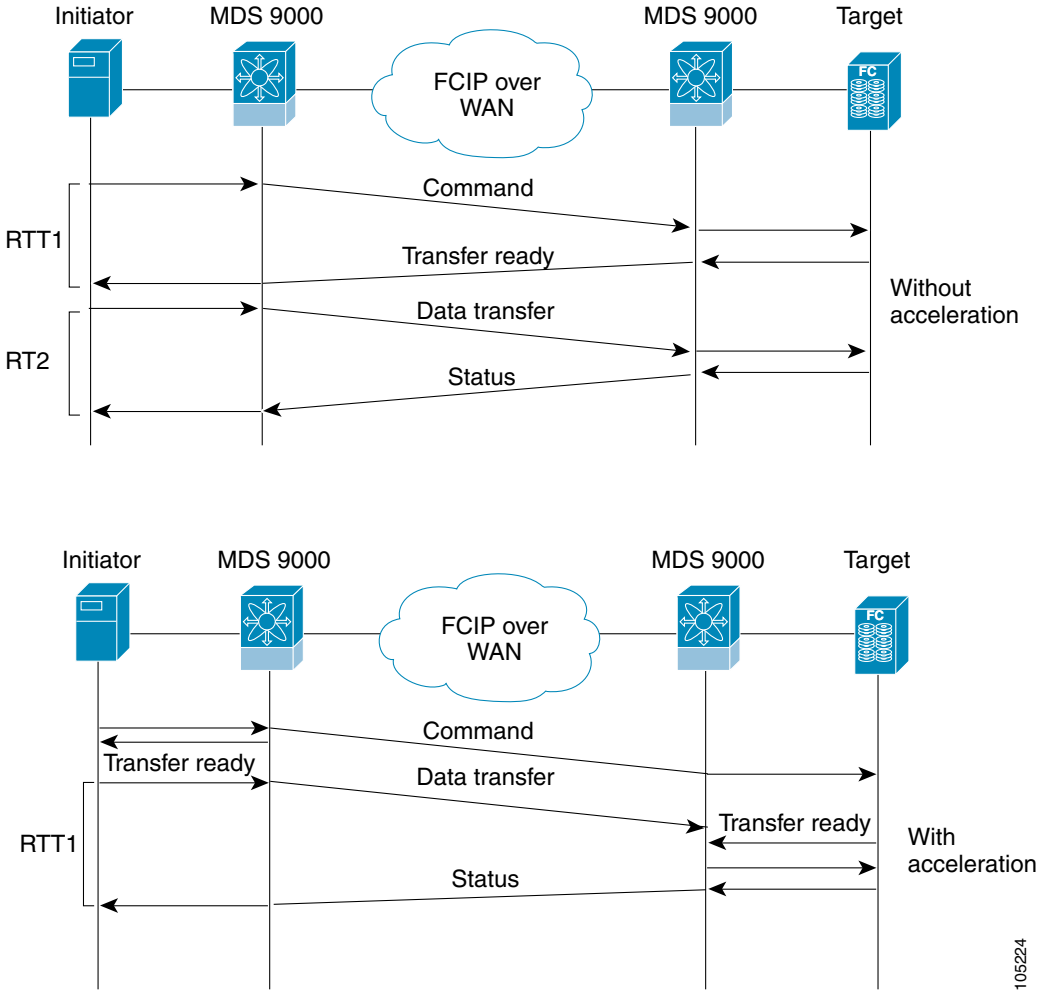

Tip FCIP write acceleration can be enabled for multiple FCIP tunnels if the tunnels are part of a dynamic PortChannel. FCIP write acceleration does not work if multiple non-PortChannel ISLs exist with equal weight between the initiator and the target port. Such a configuration might cause either SCSI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

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

Caution FCIP write acceleration with FCIP ports as members of PortChannels in Cisco MDS SAN-OS Release 2.0(1b) and later are incompatible with the FCIP write acceleration in earlier releases.

Note Write acceleration cannot be used across FSPF equal cost paths in FCIP deployments. Native Fibre Channel write acceleration can be used with Port Channels. Also, FCIP write acceleration can be used in Port Channels constructed with Port Channel Protocol (PCP).

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

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

Step 1 Choose **ISLs > FCIP** from the Physical Attributes pane on Fabric Manager. You see the FCIP profiles and links in the Information pane.
On Device manager, choose **IP > FCIP**. You see the FCIP dialog box.

Step 2 Click the **Tunnels** tab. You see the FICP link information.

Step 3 Check or uncheck the **WriteAccelerator** check box.

Step 4 Click the **IP Compression radio** button for the appropriate compression ratio in the dialog box.

Step 5 Click **Apply Changes** icon to save these changes or click **Undo Changes** to discard any unsaved changes.

FCIP Tape Acceleration

Tapes are storage devices that store and retrieve user data sequentially. Cisco MDS SAN-OS provides both tape write and read acceleration.

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

In an example of tape acceleration for write operations, the backup server in Figure 43-14 issues write operations to a drive in the tape library. Acting as a proxy for the remote tape drives, the local Cisco MDS switch proxies a transfer ready to signal the host to start sending data. After receiving all the data, the local Cisco MDS switch proxies the successful completion of the SCSI WRITE operation. This response allows the host to start the next SCSI WRITE operation. This proxy method results in more data being sent over the FCIP tunnel in the same time period compared to the time taken to send data without proxying. The proxy method improves the performance on WAN links.

Figure 43-18 FCIP Link Tape Acceleration for Write Operations

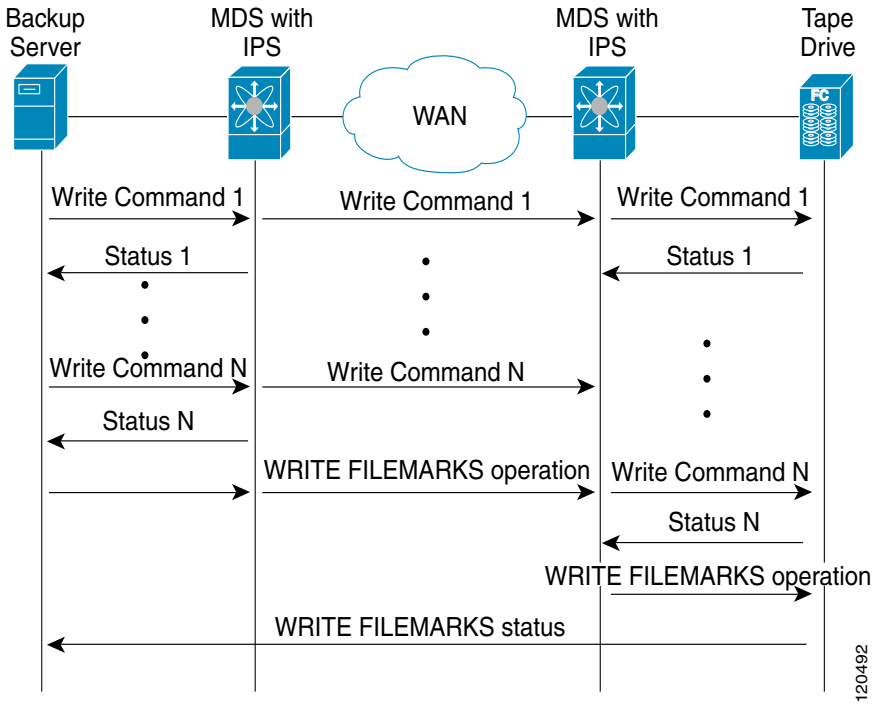

At the tape end of the FCIP tunnel, another Cisco MDS switch buffers the WRITE and data it has received. It then acts as a backup server to the tape drive by listening to a transfer ready from the tape drive before forwarding the data.

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

In an example of tape acceleration for read operations, the restore server in Figure 43-19 issues read operations to a drive in the tape library. During the restore process, the remote Cisco MDS switch at the tape end, in anticipation of more SCSI READ operations from the host, sends out SCSI READ operations on its own to the tape drive. The prefetched read data is cached at the local Cisco MDS switch. The local Cisco MDS switch on receiving SCSI READ operations from the host, sends out the cached data. This method results in more data being sent over the FCIP tunnel in the same time period compared to the time taken to send data without read acceleration for tapes. This improves the performance for tape reads on WAN links.

Figure 43-19 FCIP Link Tape Acceleration for Read Operations

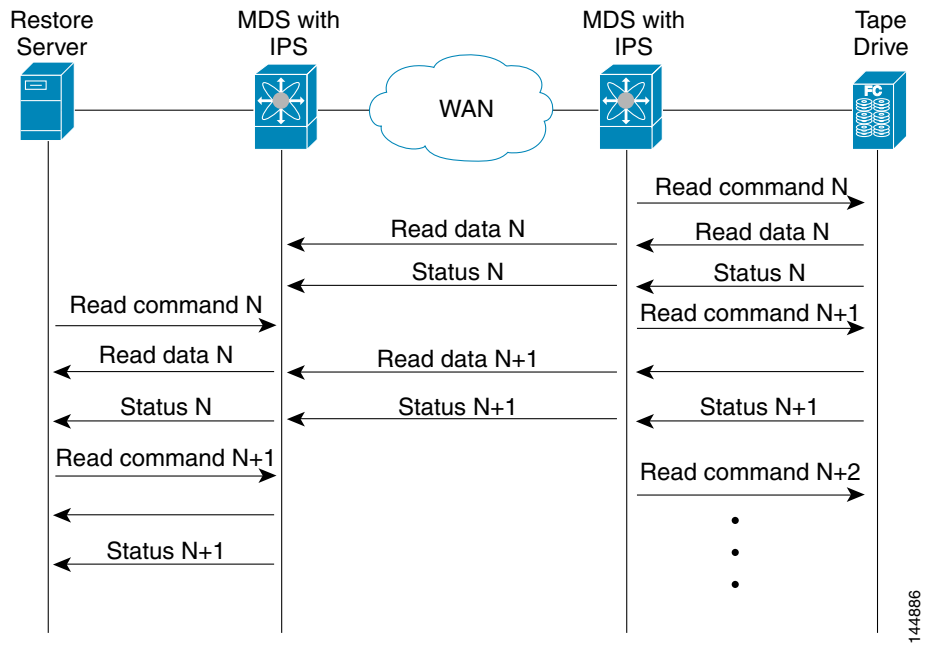

The Cisco SAN-OS provides reliable data delivery to the restore application using TCP/IP over the WAN. While tape media errors during the read operation are returned to the restore server for error handling, the Cisco SAN-OS software recovers from any other errors.

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

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

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

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

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

are flow controlled by the remote Cisco MDS switch, by not issuing any further reads. On completion of a read operation, when some data buffers are freed, the remote Cisco MDS switch resumes issuing reads.

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

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

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

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

To enable tape acceleration, follow these steps:

Step 1 From Fabric Manager, choose **ISLs > FCIP** from the Physical Attributes pane. You see the FCIP profiles and links in the Information pane.

From Device Manager, choose **IP > FCIP**. You see the FCIP dialog box.

Step 2 Click the **Tunnels** tab. You see the FICP link information.

Step 3 Click the **Create Row** icon in Fabric Manager or the **Create** button in Device Manager. You see the FCIP Tunnels dialog box.

Step 4 Set the profile ID in the ProfileID field and the tunnel ID in the TunnelID fields.

Step 5 Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.

Step 6 Check the **TapeAccelerator** check box.

Step 7 Optionally set the other fields in this dialog box and click **Create** to create this FCIP link.

Tape Library LUN Mapping for FCIP Tape Acceleration

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

Figure 43-20 shows tape drives connected to Switch 2 through a single target port. If the tape library provides LUN mapping, then all the four tape drives should be assign unique LUNs.

Figure 43-20 FCIP LUN Mapping Example

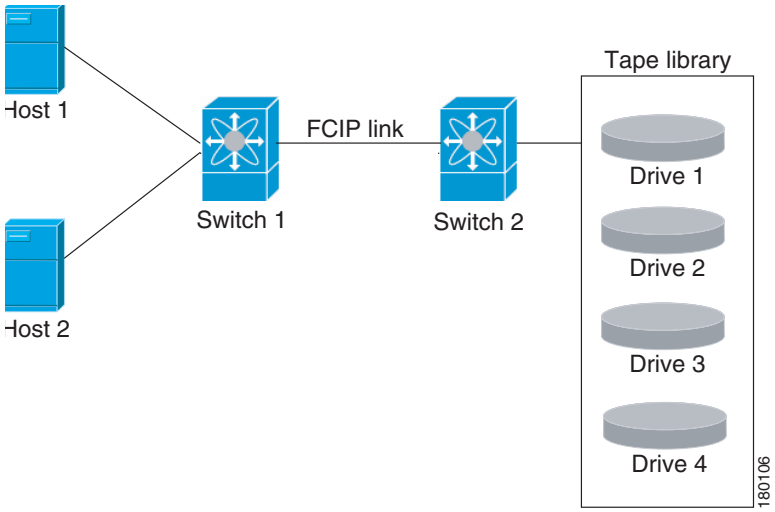

For the mappings described in Table 43-1 and Table 43-2, Host 1 has access to Drive 1 and Drive 2, and Host 2 has access to Drive 3 and Drive 4.

Table 43-1 describes correct tape library LUN mapping.

Table 43-1 Correct LUN Mapping Example With Single Host Access

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

Table 43-2 describes incorrect tape library LUN mapping.

Table 43-2 Incorrect LUN Mapping Example With Single Hosts Access

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

Another example setup is when a tape drive is shared by multiple hosts through a single tape port. For instance, Host 1 has access to Drive1 and Drive2, and Host 2 has access to Drive 2, Drive 3, and Drive 4. A correct LUN mapping configuration for such a setup is shown in Table 43-3.

Table 43-3 Correct LUN Mapping Example With Multiple Host Access

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

FCIP Compression

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

You can configure FCIP compression using one of the following modes:

- **mode1** is a fast compression mode for high bandwidth links (> 25 Mbps).
- **mode2** is a moderate compression mode for moderately low bandwidth links (between 10 and 25 Mbps).
- **mode3** is a high compression mode for low bandwidth links (< 10 Mbps).
- **auto** (default) mode picks the appropriate compression scheme based on the bandwidth of the link (the bandwidth of the link configured in the FCIP profile's TCP parameters).

The IP compression feature behavior differs between the IPS module and the MPS-14/2 module—while **mode2** and **mode3** perform software compression in both modules, **mode1** performs hardware-based compression in MPS-14/2 modules and software compression in IPS modules.

Note The Cisco MDS 9216i Switch also supports the IP compression feature. The integrated supervisor module has the same hardware components that are available in the MPS-14/2 module.

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

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

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

Default Settings

Table 43-4 lists the default settings for FCIP parameters.

Table 43-4 Default FCIP Parameters

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

CHAPTER 44

Configuring the SAN Extension Tuner

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

This chapter includes the following sections:

About the SAN Extension Tuner

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

- The TCP parameters for the FCIP profile (see the "Window Management" section on page 43-15).
- The number of concurrent SCSI I/Os generated by the application.
- The transfer size used by the application over an FCIP link.

SET is implemented in IPS ports. When enabled, this feature can be used to generate SCSI I/O commands (read and write) to the virtual target based on your configured options (see Figure 44-1).

Figure 44-1 SCSI Command Generation to the Virtual Target

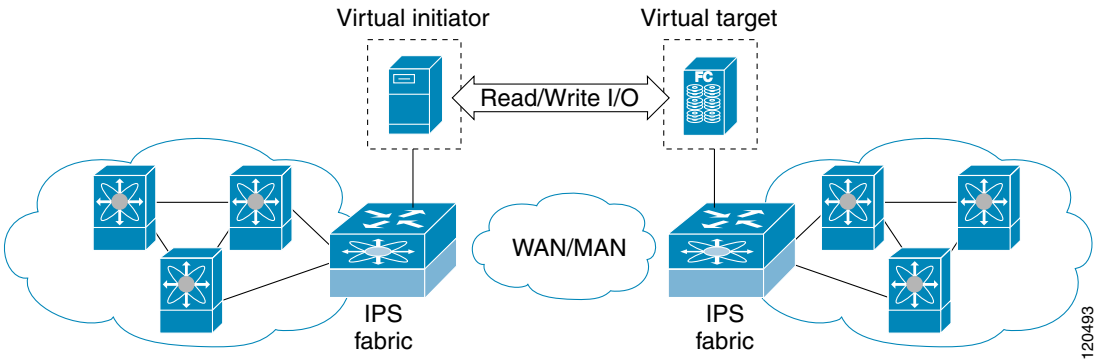

The SET feature assists with tuning by generating varying SCSI traffic workloads. It also measures throughput and response time per I/ O over an FCIP link.

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

- Following these implementation details:
  - The tuned configuration is not persistent.
  - The virtual N ports created do not register FC4 features supported with the name server. This is to avoid the hosts in the SAN from discovering these N ports as regular initiators or targets.
  - Login requests from other initiators in the SAN are rejected.
  - The virtual N ports do not implement the entire SCSI suite; it only implements the SCSI read and write commands.
  - Tuner initiators can only communicate with tuner targets.
- Verify that the Gigabit Ethernet interface is up at the physical layer (GBIC and Cable connected—an IP address is not required).
- Enable iSCSI on the switch (no other iSCSI configuration is required).
- Enable the interface (no other iSCSI interface configuration is required) (see the "Creating iSCSI Interfaces" section on page 45-5).
- Configure the virtual N ports in a separate VSAN or zone as required by your network.

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

SAN Extension Tuner Setup

Figure 44-2 provides a sample physical setup in which the virtual N ports are created on ports that are not a part of the FCIP link for which the throughput and latency is measured.

Figure 44-2 N Port Tuning Configuration Physical Example

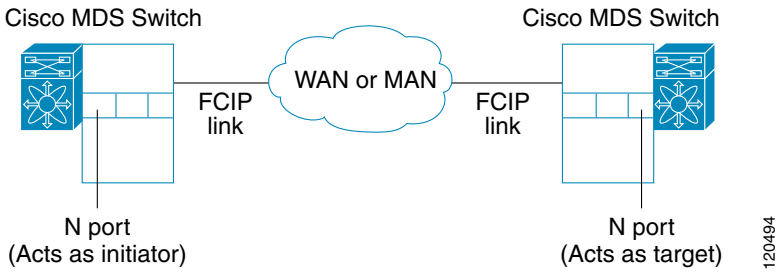

Figure 44-3 provides a sample logical setup in which the virtual N ports are created on ports that are not a part of the FCIP link for which the throughput and latency is measured.

Figure 44-3 Logical Example of N Port Tuning for a FCIP Link

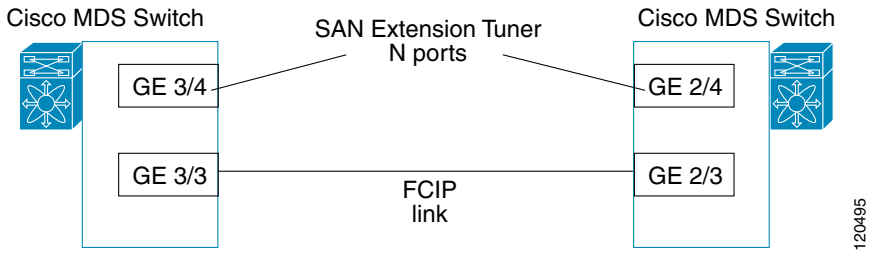

Data Pattern

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

License Prerequisites

To use the SET, you need to obtain the SAN_EXTN_OVER_IP license (see Chapter 10, "Obtaining and Installing Licenses").

Configuring the SAN Extension Tuner

This section includes the following topics:

Tuning Guidelines

To tune the required FCIP link, follow these steps:

Step 1 Configure the nWWN for the virtual N ports on the switch.

Step 2 Enable iSCSI on the interfaces on which you want to create the N ports.

Step 3 Configure the virtual N port on either side of the FCIP link.

Step 4 Ensure that the virtual N ports are not visible to real initiators in the SAN. You can use zoning (see Chapter 26, "Configuring and Managing Zones") or VSANs (see Chapter 23, "Configuring and Managing VSANs") to segregate the real initiators. Ensure that the zoning configuration is setup to allow the virtual N-ports to communicate with each other.

Step 5 Start the SCSI read and write I/Os.

Step 6 Add more N ports (as required) to other Gigabit Ethernet ports in the switch to obtain maximum throughput. One scenario that may require additional N ports is if you use FCIP PortChannels.

Using the SAN Extension Tuner Wizard

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

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

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

Step 1 Select and right-click a valid FC IP link in the Map pane then select **SAN Extension Tuner**. Or, highlight the link and choose **Tools > Other > SAN Extension Tuner**. You see the SAN Extension Tuner Wizard.

Step 2 Select the Ethernet port pairs that correspond to the FCIP link you want to tune and click **Next**.

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

Step 3 Create and activate a new zone to ensure that the virtual N ports are not visible to real initiators in the SAN by clicking **Yes** to the zone creation dialog box.

Step 4 Optionally, change the default settings for the transfer data size and the number of concurrent SCSI read and write commands as follows:

a. Set Transfer Size to the number of bytes that you expect your applications to use over the FCIP link.

b. Set Read I/0 to the number of concurrent SCSI read commands you expect your applications to generate over the FCIP link.

c. Set Write I/0 to the number of concurrent outstanding SCSI write commands you expect your applications to generate over the FCIP link.

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

d. Check the **Use Pattern File** check box and select a file that you want to use to set the data pattern that is generated by the SAN extension tuner. See the "Data Pattern" section on page 44-3.

Step 5 Click **Next**.

Step 6 Click **Start** to start the tuner. The tuner sends a continuous stream of traffic until you select **Stop**.

Step 7 Click **Show** to see the latest tuning statistics. You can select this while the tuner is running or after you stop it.

Step 8 Click **Stop** to stop the SAN extension tuner.

Default Settings

Table 44-1 lists the default settings for tuning parameters.

Table 44-1 Default Tuning Parameters

Parameters	Default
Tuning	Disabled.
Transfer ready size	Same as the transfer size in the SCSI **write** command.
Outstanding I/Os	1.
Number of transactions	1.

Table 44-1 Default Tuning Parameters (continued)

Parameters	Default
Data generation format	All-zero format.
File mark frequency	0

CHAPTER 45

Configuring iSCSI

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

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

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

Note For information on configuring Gigabit Ethernet interfaces, see the "Configuring Gigabit Ethernet Interfaces for IPv4" section on page 47-4.

This chapter includes the following sections:

About iSCSI

The iSCSI feature consists of routing iSCSI requests and responses between iSCSI hosts in an IP network and Fibre Channel storage devices in the Fibre Channel SAN that are accessible from any Fibre Channel interface of the Cisco MDS 9000 Family switch (see Figure 45-1).

Figure 45-1 Transporting iSCSI Requests and Responses for Transparent iSCSI Routing

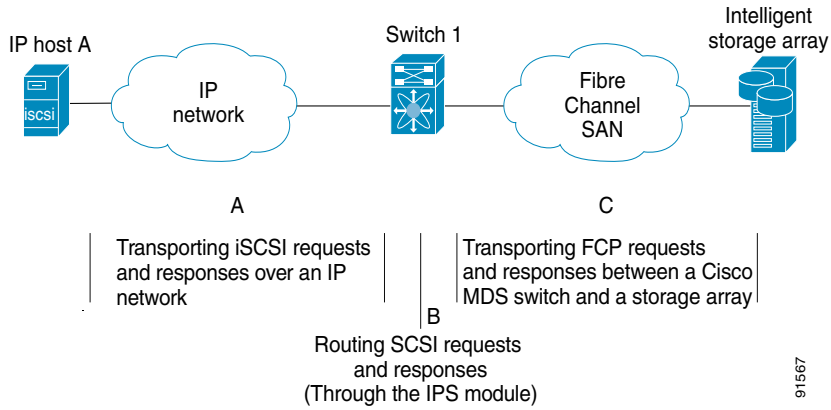

Each iSCSI host that requires access to storage through the IPS module or MPS-14/2 module needs to have a compatible iSCSI driver installed. (The Cisco.com website at http://www.cisco.com/cgi-bin/tablebuild.pl/sn5420-scsi provides a list of compatible drivers.) Using the iSCSI protocol, the iSCSI driver allows an iSCSI host to transport SCSI requests and responses over an IP network. From the host operating system perspective, the iSCSI driver appears to be a SCSI transport driver similar to a Fibre Channel driver in the host.

The IPS module or MPS-14/2 module provides transparent SCSI routing. IP hosts using the iSCSI protocol can transparently access targets on the Fibre Channel network. Figure 45-1 provides an example of a typical configuration of an iSCSI host connected to an IPS module or MPS-14/2 module through the IP network to access Fibre Channel storage on the Fibre Channel SAN.

The IPS module or MPS-14/2 module create a separate iSCSI SAN view and Fibre Channel SAN view. For the iSCSI SAN view, the IPS module or MPS-14/2 module creates iSCSI virtual targets and then maps them to physical Fibre Channel targets available in the Fibre Channel SAN. They present the Fibre Channel targets to IP hosts as if the physical iSCSI targets were attached to the IP network (see Figure 45-2).

Figure 45-2 iSCSI SAN View—iSCSI Virtual Targets

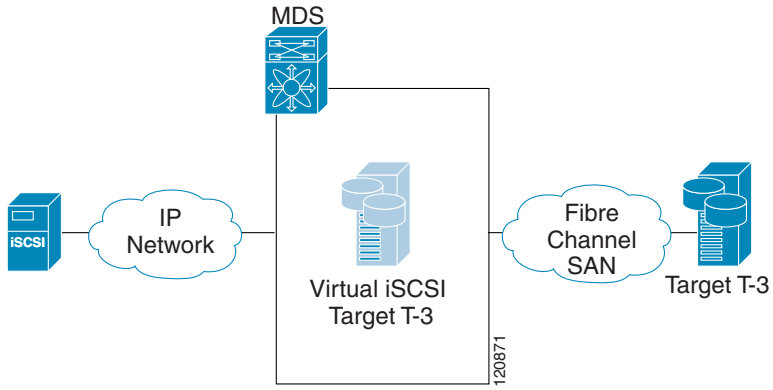

For the Fibre Channel SAN view, the IPS module or MPS-14/2 module presents iSCSI hosts as a virtual Fibre Channel host. The storage devices communicate with the virtual Fibre Channel host similar to the way communications are performed with real Fibre Channel hosts (see Figure 45-3).

Figure 45-3 Fibre Channel SAN View—iSCSI Host as an HBA

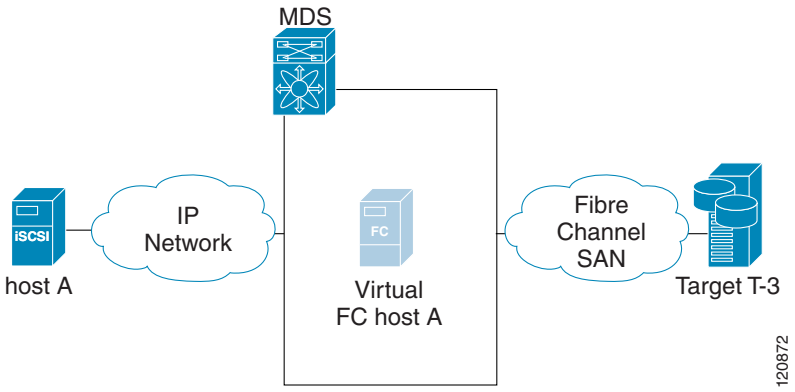

The IPS modules or MPS-14/2 modules transparently map between the iSCSI virtual target and the virtual Fibre Channel host (see Figure 45-4).

Figure 45-4 iSCSI to FCP (Fibre Channel) Routing

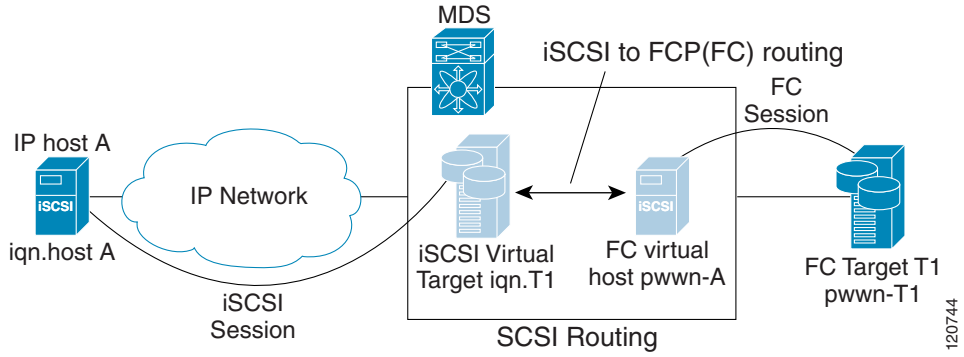

Routing SCSI from the IP host to the Fibre Channel storage device consists of the following main actions:

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

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

Configuring iSCSI

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

Enabling iSCSI

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

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

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

You see the iSCSI tables in the Information pane shown in Figure 45-5.

Figure 45-5 iSCSI Tables in Fabric Manager

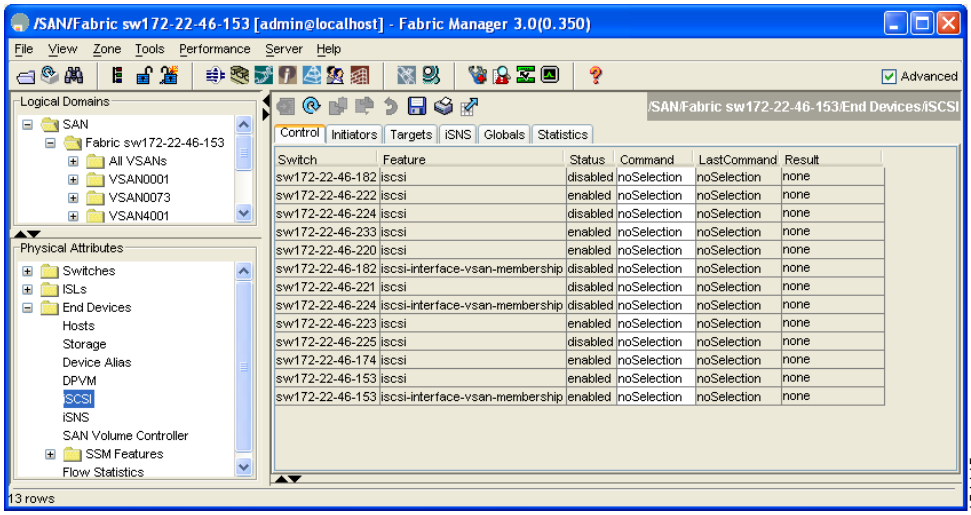

The **Control** tab is the default tab. You see the iSCSI enable status for all switches in the fabric that contain IPS ports.

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

Step 3 Click the **Apply Changes** icon to save these changes or click the **Undo Changes** icon to remove all changes.

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

Creating iSCSI Interfaces

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

Using the iSCSI Wizard

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

Step 1 Click the **iSCSI Setup Wizard** icon.

Step 2 Select an existing iSCSI initiator or add the iSCSI node name or IP address for a new iSCSI initiator.

Step 3 Select the switch for this iSCSI initiator if you are adding a new iSCSI initiator and click **Next**.

Step 4 Select the VSAN and targets to associate with this iSCSI initiator (see Figure 45-6) and click **Next**.

Figure 45-6 Select Targets

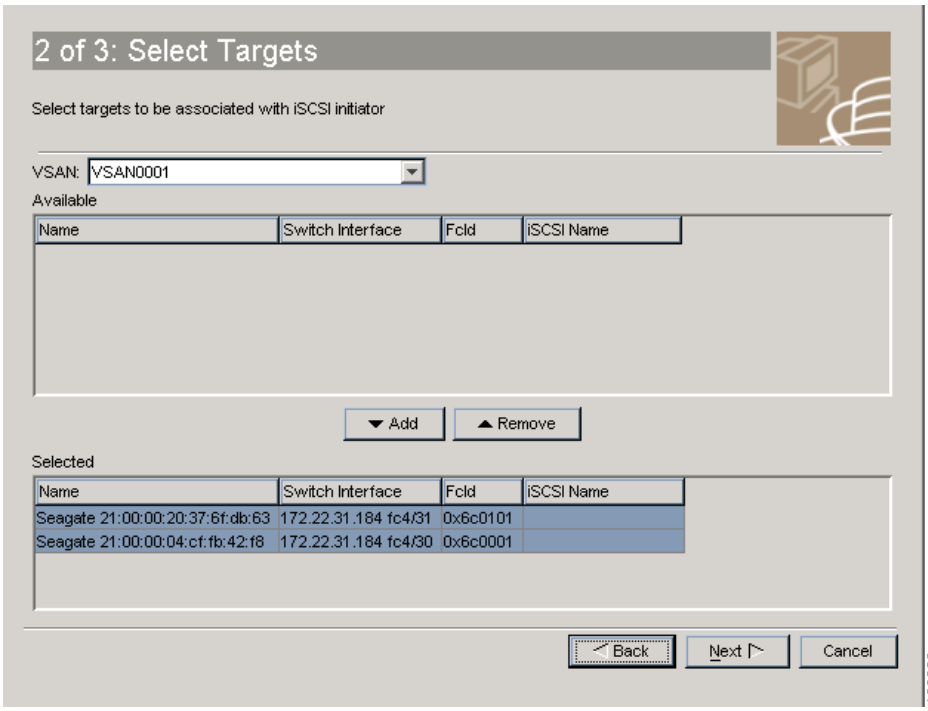

Step 5 Set the zone name for this new iSCSI zone and optionally check the **Read Only** check box (see Figure 45-7).

Figure 45-7 Select Zone

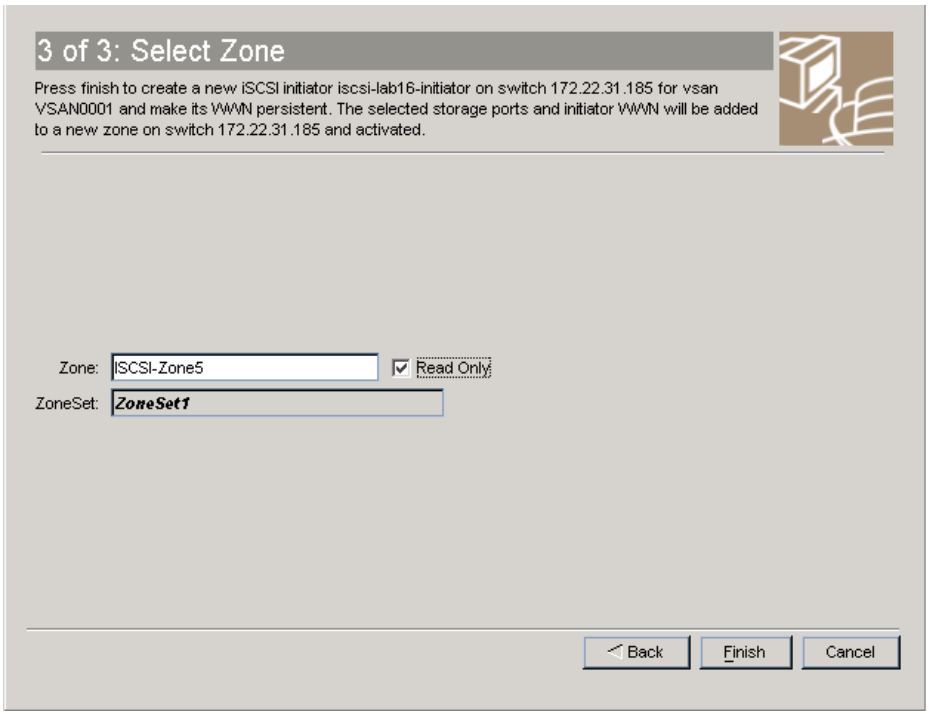

Step 6 Click **Finish** to create this iSCSI initiator or click **Cancel** to close the wizard without creating the iSCSI initiator. If created, the target VSAN is added to the iSCSI host VSAN list.

Presenting Fibre Channel Targets as iSCSI Targets

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

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

  Static mapping should be used when iSCSI hosts should be restricted to subsets of LUs in the Fibre Channel targets and/or iSCSI access control is needed (see the "iSCSI Access Control" section on page 45-27). Also, static mapping allows the configuration of transparent failover if the LUs of the Fibre Channel targets are reachable by redundant Fibre Channel ports (see the "Transparent Target Failover" section on page 45-58).

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

Dynamic Mapping

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

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

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

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

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

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

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

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

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

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

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

For example, if an iSCSI target was created for a Fibre Channel target port with pWWN 31:00:11:22:33:44:55:66 and that pWWN contains LUN 0, LUN 1, and LUN 2, those LUNs would become available to an IP host through the iSCSI target node name iqn.1987-05.com.cisco:05.*MDS_switch_management_IP_address* 01-01.3100112233445566 (see Figure 45-8).

Figure 45-8 Dynamic Target Mapping

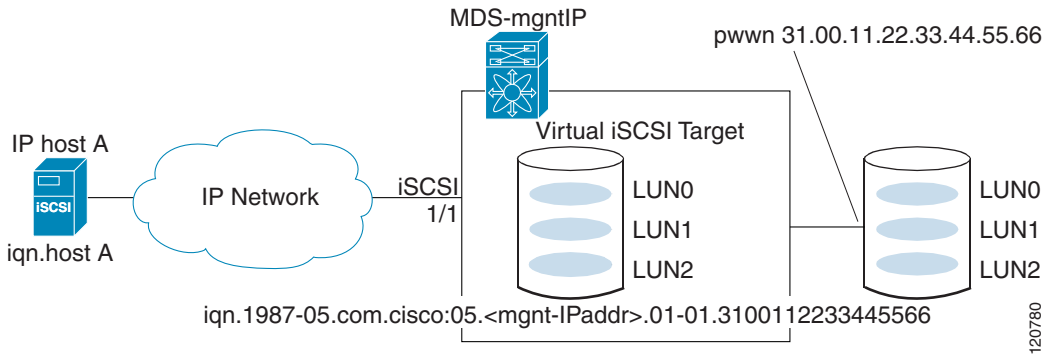

Note Each iSCSI initiator may not have access to all targets depending on the configured access control mechanisms (see the "iSCSI Access Control" section on page 45-27).

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

Step 1 Click **IP > iSCSI**.

You see the iSCSI tables under the Initiator tab.

Figure 45-9 iSCSI Tables in Device Manager

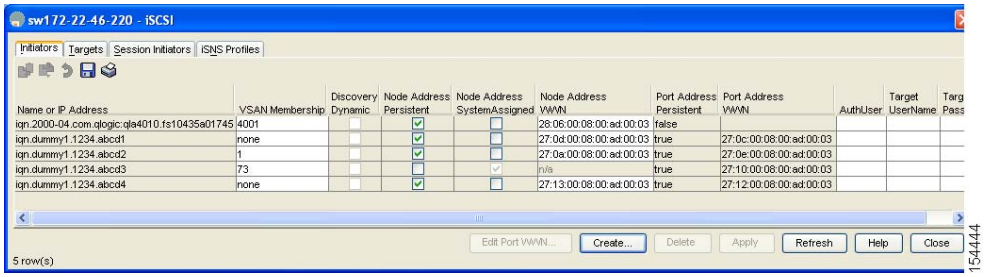

Step 2 Click the **Targets** tab to display a list of existing iSCSI targets shown in Figure 45-10.

Figure 45-10 iSCSI Targets in Device Manager

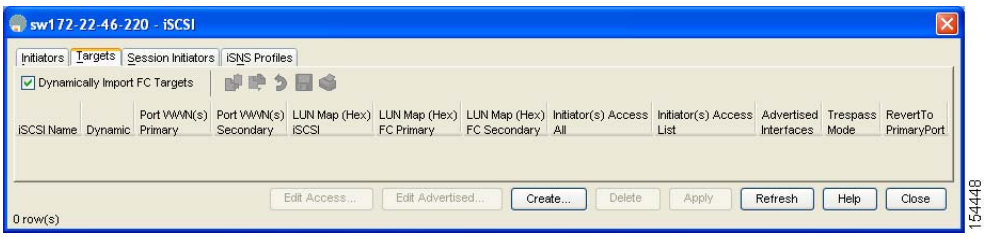

Step 3 Check the **Dynamically Import FC Targets** check box.

Step 4 Click **Apply** to save this change or close the dialog box without saving any changes.

Static Mapping

You can manually (statically) create an iSCSI target by assigning a user-defined unique iSCSI node name to it. The iSCSI qualified name is restricted to a minimum length of 16 characters and a maximum of 223 characters. A statically mapped iSCSI target can either map the whole Fibre Channel target port (all LUNs in the target port mapped to the iSCSI target), or it can contain one or more LUs from a Fibre Channel target port (see Figure 45-11).

Figure 45-11 Statically Mapped iSCSI Targets

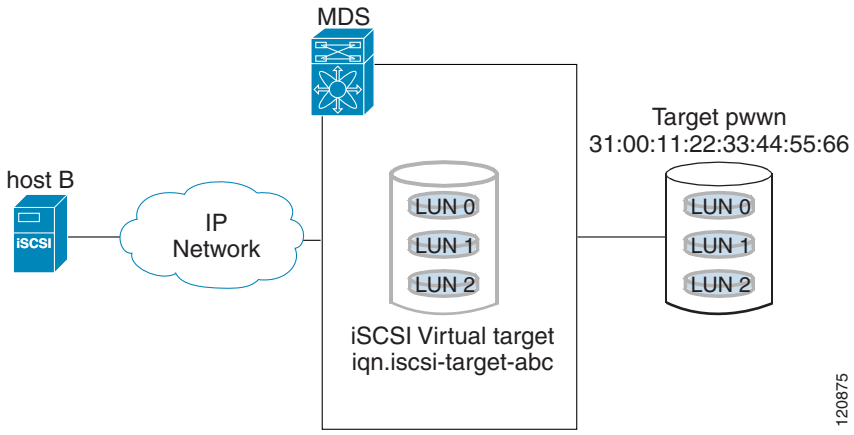

To create a static iSCSI virtual target for the entire Fibre Channel target port using Device Manager, follow these steps:

Step 1 Click **IP > iSCSI**.

You see the iSCSI tables under the Initiator tab.

Figure 45-12 iSCSI Tables in Device Manager

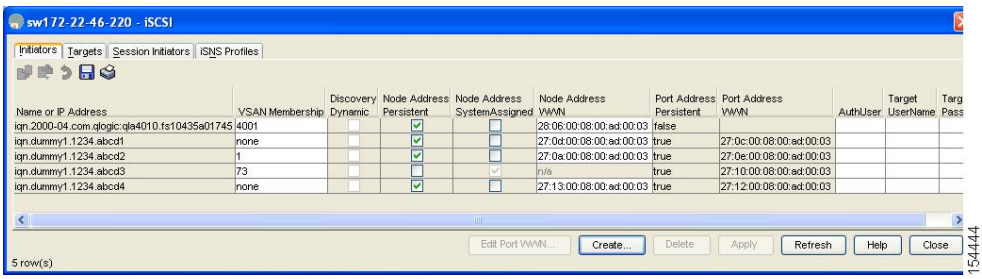

Step 2 Click the **Targets** tab to display a list of existing iSCSI targets shown in Figure 45-13.

Figure 45-13 iSCSI Targets in Device Manager

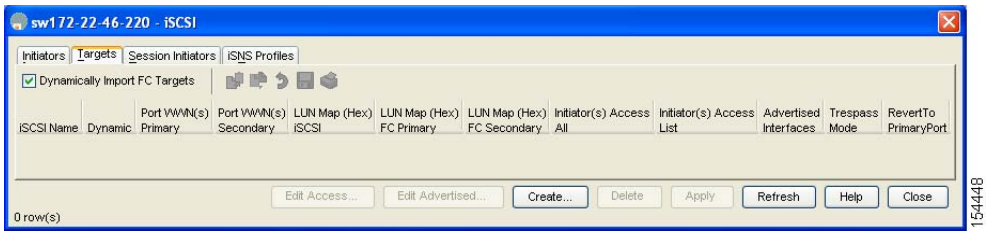

Step 3 Click **Create** to create an iSCSI target.

You see the Create iSCSI Targets dialog box shown in Figure 45-14.

Figure 45-14 Create iSCSI Targets

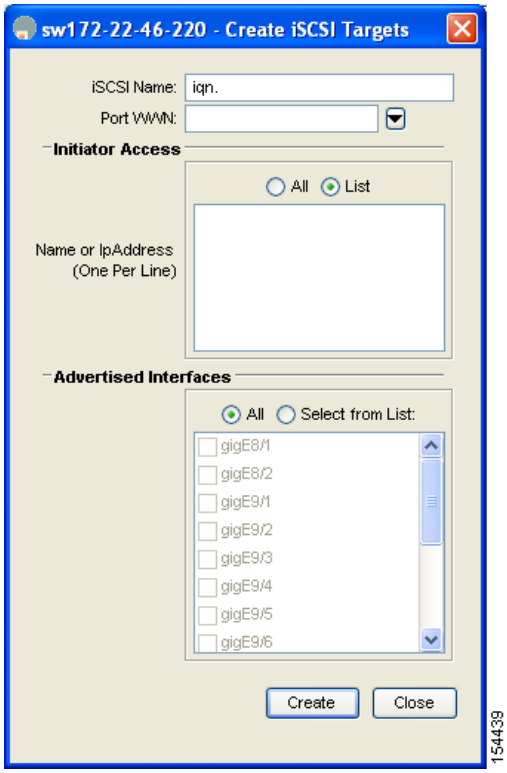

Step 4 Set the iSCSI target node name in the iSCSI Name field, in IQN format.

Step 5 Set the Port WWN field for the Fibre Channel target port you are mapping.

Step 6 Choose the **List** radio button and set the iSCSI initiator node names or IP addresses that you want this virtual iSCSI target to access, or choose the **All** radio button to let the iSCSI target access all iSCSI initiators (see Figure 45-14). Also see the "iSCSI Access Control" section on page 45-27.

Step 7 Choose the **Selected from List** radio button and check each interface you want to advertise the iSCSI targets on or choose the **All** radio button to advertise all interfaces.

Step 8 Click **Apply** to save this change or click **Cancel** to close the dialog box without saving any changes.

Tip An iSCSI target cannot contain more than one Fibre Channel target port. If you have already mapped the whole Fibre Channel target port, you cannot use the LUN mapping option.

Note See the "iSCSI-Based Access Control" section on page 45-29 for more information on controlling access to statically mapped targets.

Advertising Static iSCSI Targets

You can limit the Gigabit Ethernet interfaces through which static iSCSI targets are advertised. By default iSCSI targets are advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces.

To configure a specific interface that should advertise the iSCSI virtual target using Device Manager, follow these steps:

Step 1 Click **IP > iSCSI**.

You see the iSCSI Configuration.

Figure 45-15 iSCSI Tables in Device Manager

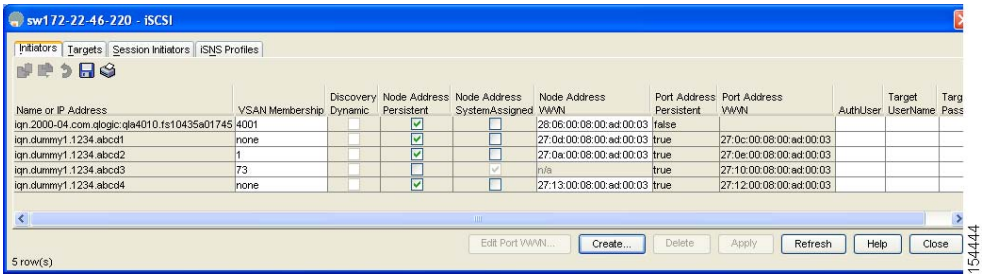

Step 2 Click the **Targets** tab to display a list of existing iSCSI targets

Step 3 Right-Click the iSCSI target that you want to modify and click **Edit Advertised**.

You see the Advertised Interfaces dialog box.

Step 4 Optionally, right-click on an interface that you want to delete and select **Delete**.

Step 5 Optionally, click **Create** to advertise on more interfaces.

You see the Create Advertised Interfaces dialog box.

Step 6 Select an interface from the Interfaces drop-down menu and then click **Create** to add this to the list of advertised interfaces, or click **Close** to close the dialog box without saving any changes.

iSCSI Virtual Target Configuration Examples

This section provides three examples of iSCSI virtual target configurations.

Example 1

This example assigns the whole Fibre Channel target as a iSCSI virtual target. All LUNs that are part of the Fibre Channel target are available as part of the iSCSI target (see Figure 45-16).

Figure 45-16 Assigning iSCSI Node Names

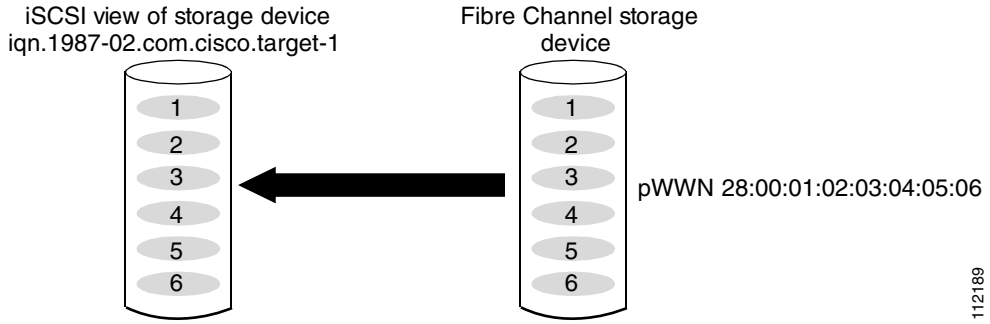


```
iscsi virtual-target name iqn.1987-02.com.cisco.target-1
    pWWN 28:00:01:02:03:04:05:06
```

Example 2

This example maps a subset of LUNs of a Fibre Channel target to three iSCSI virtual targets. Each iSCSI target only has one LUN (see Figure 45-17).

Figure 45-17 Mapping LUNs to a iSCSI Node Name

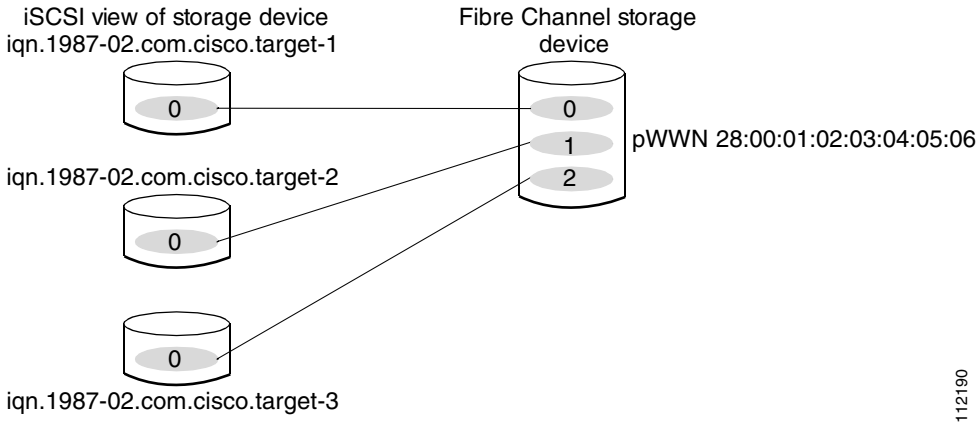


```
iscsi virtual-target name iqn.1987-02.com.cisco.target-1
    pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-2
    pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-3
```

```
    pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 0
```

Example 3

This example maps three subsets of Fibre Channel LUN targets to three iSCSI virtual targets. Two iSCSI targets have one LUN and the third iSCSI target has two LUNs (see Figure 45-18).

Figure 45-18 Mapping LUNs to Multiple iSCSI Node Names

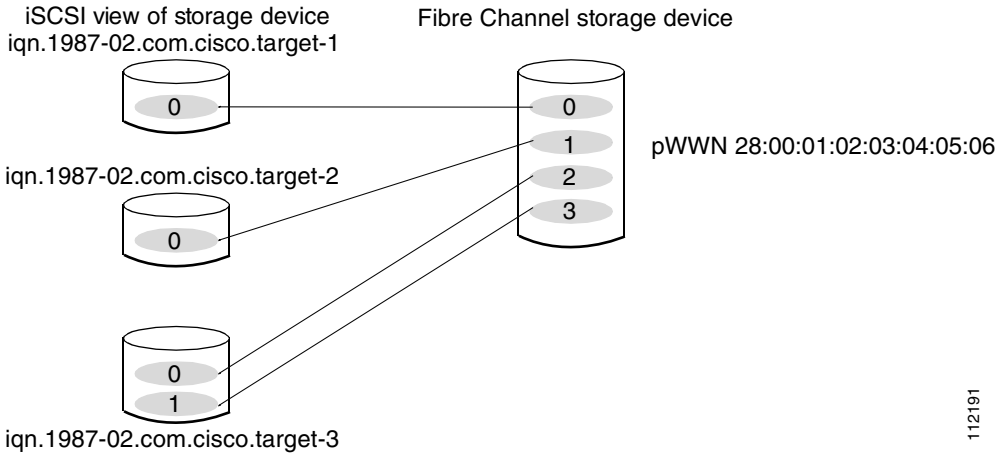


```
iscsi virtual-target name iqn.1987-02.com.cisco.target-1
    pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-2
    pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-3
    pWWN 28:00:01:02:03:04:05:06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

Presenting iSCSI Hosts as Virtual Fibre Channel Hosts

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

Initiator Identification

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

- iSCSI qualified name (IQN)

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

- IP address

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

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

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

Step 1 Choose **Interfaces > FC Logical** from the Physical Attributes pane.

You see the interfaces configuration in the Information pane.

Step 2 Select the **iSCSI** tab.

You see the iSCSI interfaces configuration.

Step 3 Right-click on the Initiator ID Mode field for the iSCSI interface that you want to modify and select **name** or **ipaddress** from the drop-down menu.

Step 4 Click **Apply Changes** to save this change or click **Undo Changes** to discard any unsaved changes.

Initiator Presentation Modes

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

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

Caution Enabling proxy initiator mode of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the “Changing iSCSI Interface Parameters and the Impact on Load Balancing” section on page 45-52.

The Cisco MDS switches support the following iSCSI session limits:

- The maximum number of iSCSI sessions on a switch is 5000.
- The maximum number of iSCSI sessions per IPS port in transparent initiator mode is 500.
- The maximum number of iSCSI sessions per IPS port in proxy initiator mode is 500.
- The maximum number of concurrent iSCSI sessions allowed per port is five.

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

Transparent Initiator Mode

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

When an iSCSI host connects to the IPS module or MPS-14/2 module, a virtual host N port (HBA port) is created for the host (see Figure 45-19). Every Fibre Channel N port requires a unique node WWN and port WWN.

Figure 45-19 Virtual Host HBA Port

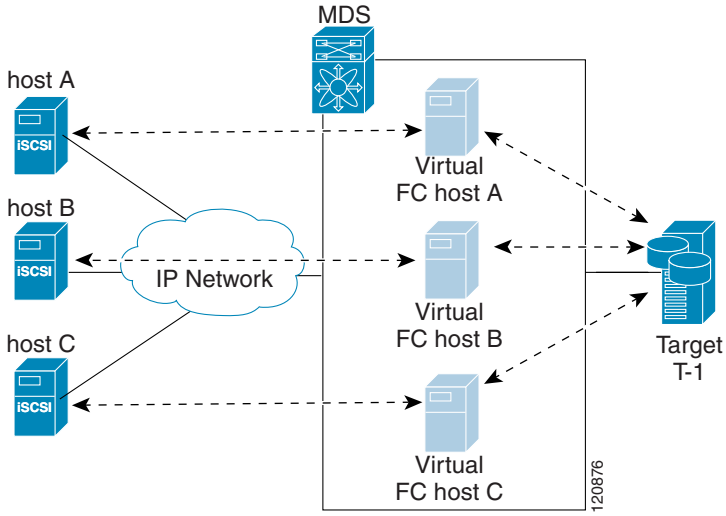

After the virtual N port is created with the WWNs, a fabric login (FLOGI) is done through the virtual iSCSI interface of the IPS port. After the FLOGI is completed, the virtual N port is online in the Fibre Channel SAN and virtual N port is registered in the Fibre Channel name server. The IPS module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

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

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

For every iSCSI session from the host to the iSCSI virtual target there is a corresponding Fibre Channel session to the real Fibre Channel target. In Figure 45-19, there are three iSCSI hosts and all three of them connect to the same Fibre Channel target. There is one Fibre Channel session from each of the three virtual Fibre Channel hosts to the target.

iSCSI Initiator Idle Timeout

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

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

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

You see the iSCSI tables in the Information pane shown in Figure 45-20.

Figure 45-20 iSCSI Tables in Fabric Manager

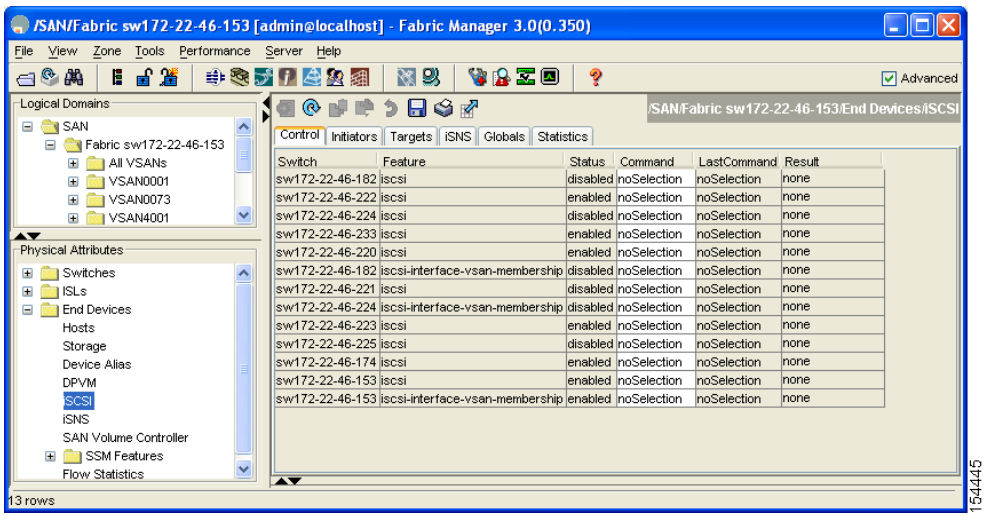

Step 2 Select the **Globals** tab.

You see the iSCSI global configuration.

Step 3 Right-click on the InitiatorIdle Timeout field that you want to modify and enter the new timeout value.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

WWN Assignment for iSCSI Initiators

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

- Dynamic mapping (default)
- Static mapping

Dynamic Mapping

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

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

The following are three dynamic initiator modes are supported:

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

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

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

Static Mapping

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

You can implement static mapping in one of two ways:

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

Tip We recommend using the system-assign option. If you manually assign a WWN, you must ensure its uniqueness (see the “World Wide Names” section on page 32-14). You should not use any previously assigned WWNs.

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

Step 1 Select **IP > iSCSI.**

You see the iSCSI configuration dialog box.

Figure 45-21 iSCSI Tables in Device Manager

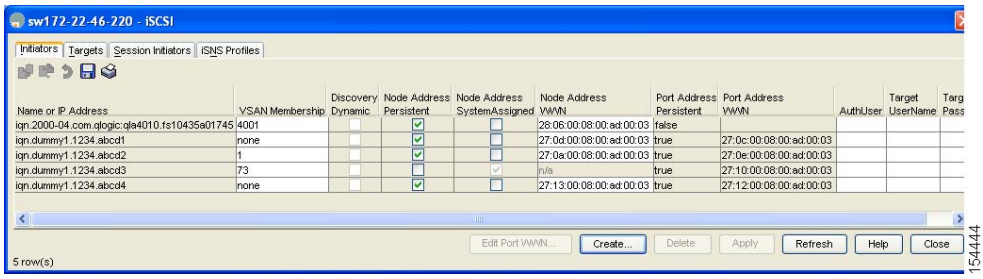

The **Initiators** tab is the default.

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

You see the Create iSCSI Initiators dialog box in Figure 45-22.

Figure 45-22 Create iSCSI Initiators Dialog Box

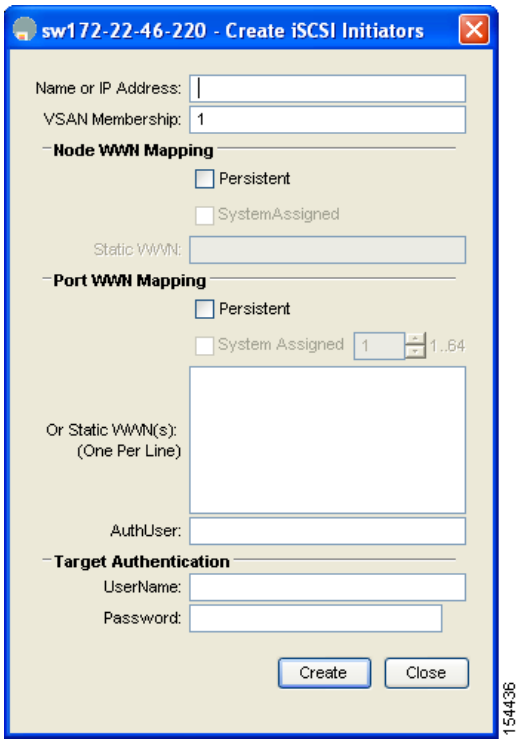

Step 3 Set the iSCSI node name or IP address and VSAN membership.

Step 4 In the Node WWN section, check the **Persistent** check box.

Step 5 Check the **System Assigned** check box if you want the switch to assign the nWWN. Or leave this unchecked and set the Static WWN field.

Step 6 In the Port WWN section, check the **Persistent** check box if you want to statically map pWWNs to the iSCSI initiator.

Step 7 If persistent, check the **System Assigned** check box and set the number of pWWNs to reserve for this iSCSI initiator if you want the switch to assign pWWNs. Alternately, you can leave this unchecked and set one or more pWWNs for this iSCSI initiator.

Step 8 Optionally set the AuthUser field if authentication is enabled. Also see the "iSCSI Session Authentication" section on page 45-31.

Step 9 Click **Create** to create this iSCSI initiator or click **Cancel** to close the dialog box without creating an iSCSI initiator.

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

Making the Dynamic iSCSI Initiator WWN Mapping Static

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

You can convert a dynamic iSCSI initiator to static iSCSI initiator and make its WWNs persistent (see "Dynamic Mapping" section on page 45-18).

Note You cannot convert a dynamic iSCSI initiator to a static iSLB initiator or a dynamic iSLB initiator to a static iSCSI initiator.

Note Making the dynamic pWWNs static after the initiator is created is supported only through the CLI, not through Device Manager or Fabric Manager. In Fabric Manager or Device Manager, you must delete and then recreate this initiator to have the pWWNs static.

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

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

You see the iSCSI tables in the Information pane shown in Figure 45-23.

Figure 45-23 iSCSI Tables in Fabric Manager

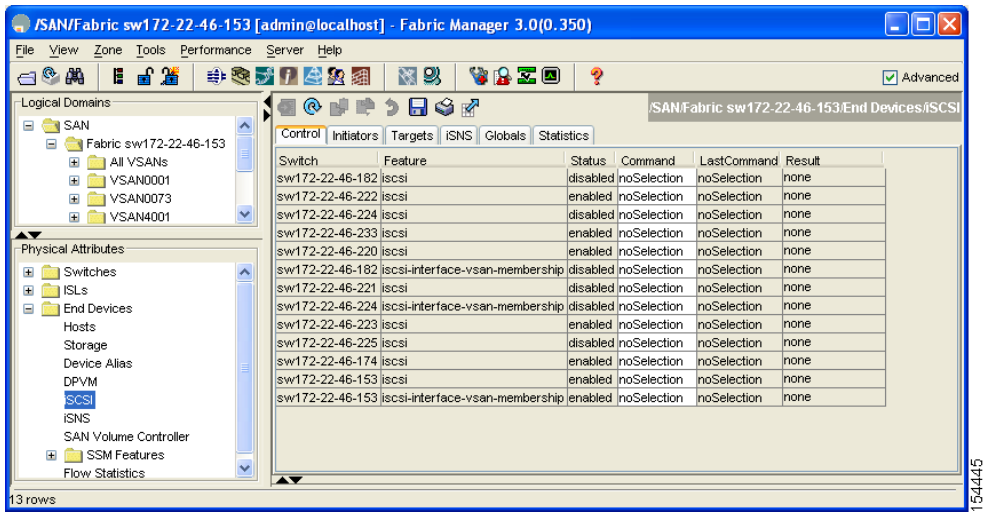

Step 2 Select the **Initiators** tab.

You see the iSCSI initiators configured.

Step 3 Check the **Persistent Node WWN** check box for the iSCSI initiators that you want to make static.

Step 4 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

Proxy-Initiator Mode

In the event that the Fibre Channel storage device requires explicit LUN access control for every host using the transparent initiator mode (presenting one iSCSI host as one Fibre Channel host) means every iSCSI host has to be configured statically. This can mean several configuration tasks for each iSCSI host. In this case, using the proxy initiator mode simplifies the configuration.

In this mode, only one virtual host N port (HBA port) is created per IPS port. All the iSCSI hosts connecting to that IPS port will be multiplexed using the same virtual host N port (see Figure 45-24). This mode simplifies the task of statically binding WWNs. LUN mapping and assignment on the Fibre Channel storage array must be configured to allow access from the proxy virtual N port's pWWN for all LUNs used by each iSCSI initiator that connects through this IPS port. The LUN is then assigned to each iSCSI initiator by configuring iSCSI virtual targets (see the "Static Mapping" section on page 45-10) with LUN mapping and iSCSI access control (see the "iSCSI Access Control" section on page 45-27).

Figure 45-24 Multiplexing IPS Ports

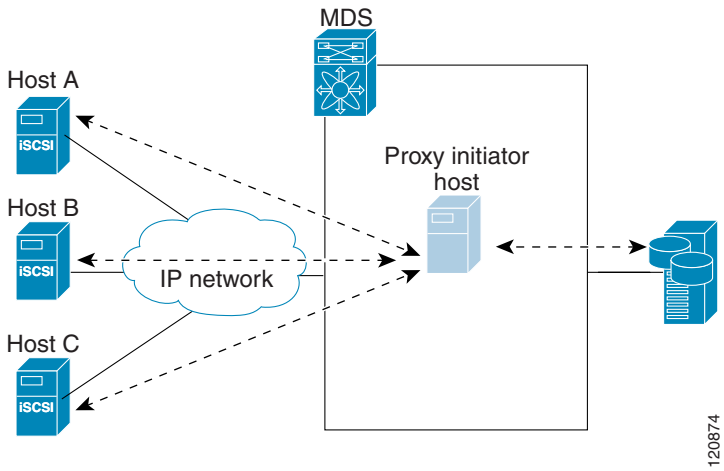

Proxy initiator mode can be configured on a per IPS port basis, in which case only iSCSI initiators terminating on that IPS port will be in this mode.

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

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

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

Caution Enabling the proxy initiator mode of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the "Changing iSCSI Interface Parameters and the Impact on Load Balancing" section on page 45-52.

To configure the proxy initiator, follow these steps:

Step 1 From Fabric Manager, expand **Switches**, expand **Interfaces** and then select **FC Logical** from the Physical Attributes pane.

You see the Interface tables in the Information pane shown in Figure 45-25.

Figure 45-25 FC Logical Interface Tables

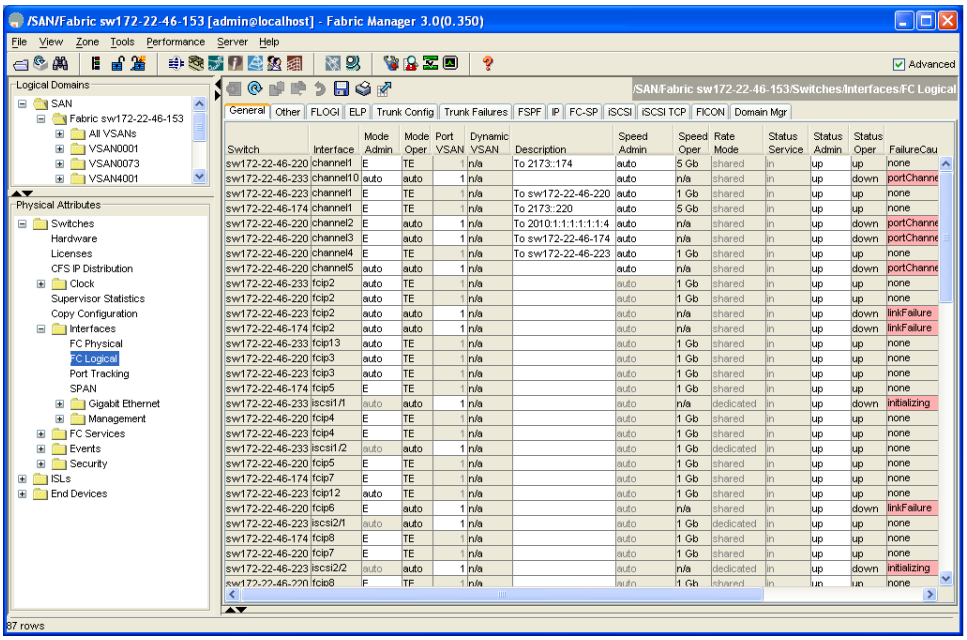

From Device Manager, click **Interfaces > Ethernet and iSCSI**.

You see the interfaces dialog box shown in Figure 45-26.

Figure 45-26 Ethernet and iSCSI Interfaces in Device e Manager

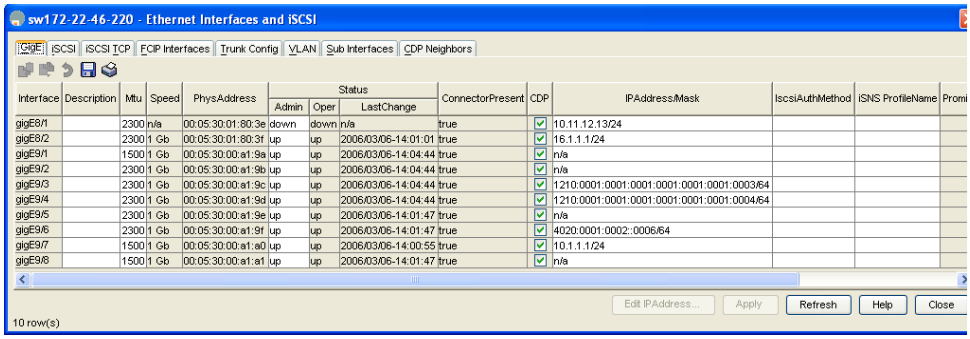

Step 2 Click the **iSCSI** tab in either FM or DM.

You see the iSCSI interface configuration table.

Figure 45-27 iSCSI Tab in Device Manager

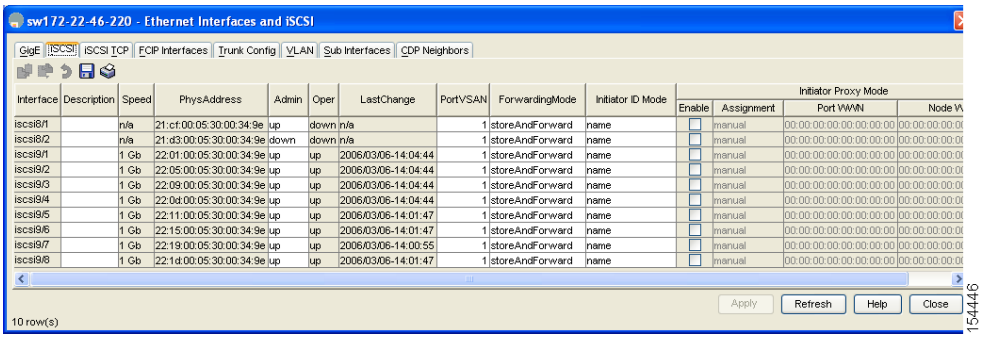

Step 3 In the Initiator Proxy Mode section, check the **Enable** check box.

Step 4 Click the **Apply Changes** icon in Fabric Manager or click **Apply** in Device Manager to save these changes. Or, click **Undo Changes** in Fabric Manager or click **Cancel** in Device Manager to discard changes.

Note When an interface is in proxy initiator mode, you can only configure Fibre Channel access control (zoning) based on the iSCSI interface's proxy N port attributes—the WWN pairs or the FC ID. You cannot configure zoning using iSCSI attributes such as IP address or IQN of the iSCSI initiator. To enforce initiator-based access control, use iSCSI based access control (see the "iSCSI Access Control" section on page 45-27).

VSAN Membership for iSCSI

Similar to Fibre Channel devices, iSCSI devices have two mechanisms by which VSAN membership can be defined.

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

VSAN Membership for iSCSI Hosts

Individual iSCSI hosts can be configured to be in a specific VSAN (similar to the DPVM feature for Fibre Channel, see Chapter 24, "Creating Dynamic VSANs"). The specified VSAN overrides the iSCSI interface VSAN membership.

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

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

You see the iSCSI tables in the Information pane shown in Figure 45-28.

Figure 45-28 iSCSI Tables in Fabric Manager

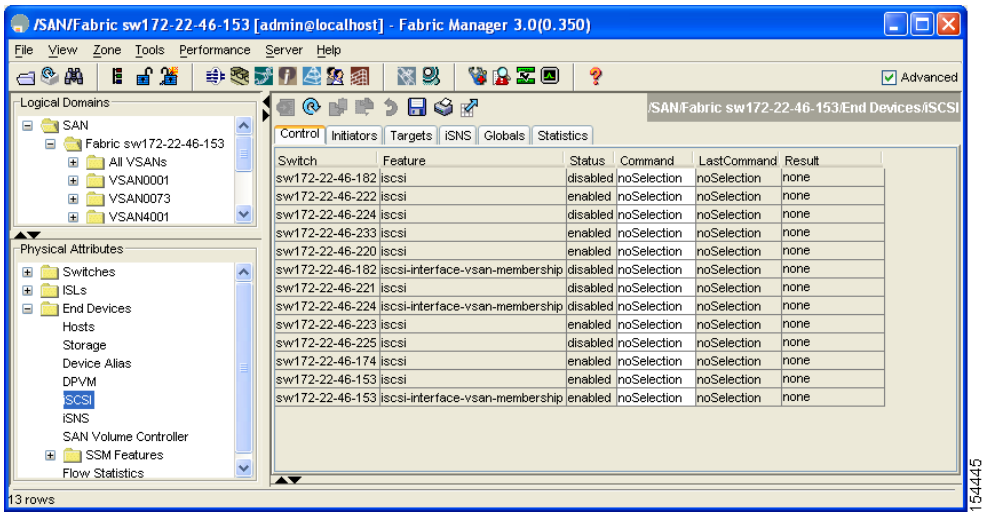

Step 2 Select the **Initiators** tab.

You see the iSCSI initiators configured.

Step 3 Fill in the VSAN Membership field to assign a VSAN to the iSCSI hosts.

Step 4 Click the **Apply Changes** icon to save these changes or click **Undo Changes** to discard any unsaved changes.

Note When an initiator is configured in any other VSAN (other than VSAN 1), for example VSAN 2, the initiator is automatically removed from VSAN 1. If you also want it to be present in VSAN 1, you must explicitly configure the initiator in VSAN 1.

VSAN Membership for iSCSI Interfaces

VSAN membership can be configured for an iSCSI interface, called the *port VSAN*. All the iSCSI devices that connect to this interface automatically become members of this VSAN, if it is not explicitly configured in a VSAN. In other words, the port VSAN of an iSCSI interface is the default VSAN for all dynamic iSCSI initiators. The default port VSAN of an iSCSI interface is VSAN 1.

Caution Changing the VSAN membership of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the "Changing iSCSI Interface Parameters and the Impact on Load Balancing" section on page 45-52.

To change the default port VSAN for an iSCSI interface using Device Manager, follow these steps:

Step 1 Click **Interfaces > Ethernet and iSCSI**.

You see the Ethernet and iSCSI interfaces dialog box.

Figure 45-29 Ethernet and iSCSI Interfaces Dialog Box in Device Manager

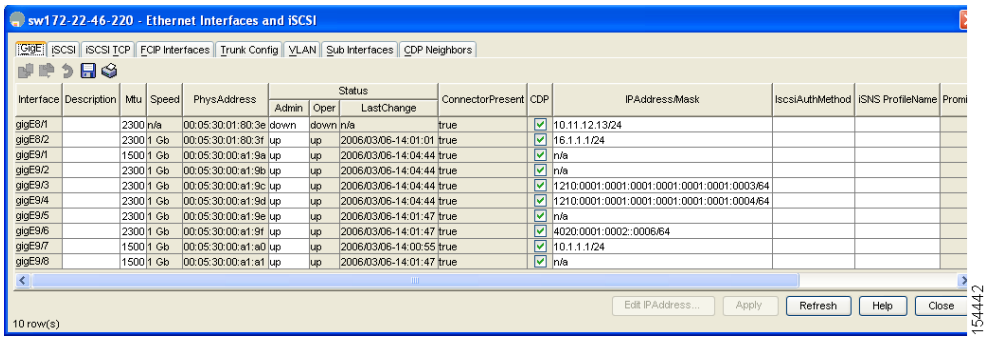

Step 2 Click the **iSCSI** tab.

You see the iSCSI interface configuration table in Figure 45-30.

Figure 45-30 iSCSI Tab in Device Manager

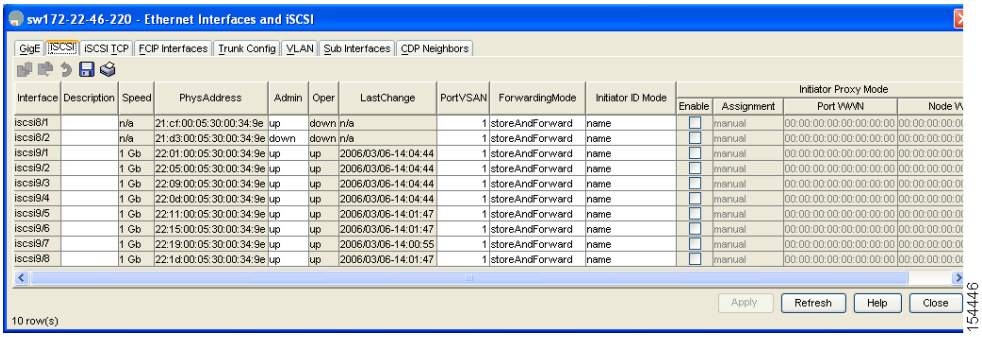

Step 3 Double-click the PortVSAN column and modify the default port VSAN.

Step 4 Click **Apply** to save these changes, or click **Cancel** to discard changes.

Example of VSAN Membership for iSCSI Devices

Figure 45-31 provides an example of VSAN membership for iSCSI devices:

- iSCSI interface 1/1 is a member of VSAN Y.
- iSCSI initiator host A has explicit VSAN membership to VSAN X.
- Three iSCSI initiators (host A, host B, and host C) connect to iSCSI interface 1/1.

Figure 45-31 VSAN Membership for iSCSI Interfaces

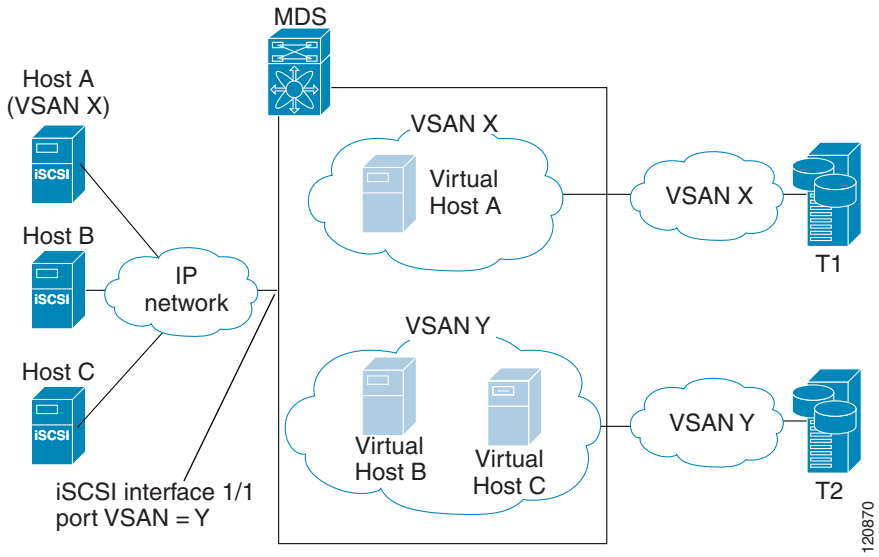

Host A's virtual Fibre Channel N port will be added to VSAN X because of explicit membership for the initiator. The virtual host-B and host-C N ports do not have any explicit membership configuration so they will inherit the iSCSI interface VSAN membership and be part of VSAN Y.

Advanced VSAN Membership for iSCSI Hosts

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

iSCSI Access Control

Two mechanisms of access control are available for iSCSI devices.

- Fibre Channel zoning-based access control
- iSCSI ACL-based access control

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

Fibre Channel Zoning-Based Access Control

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

Common mechanisms for identifying members of a Fibre Channel zone are the following (see Chapter 26, "Configuring and Managing Zones" for details on Fibre Channel zoning):

- Fibre Channel device pWWN.

- Interface and switch WWN.

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

In transparent initiator mode (where one Fibre Channel virtual N port is created for each iSCSI host as described in the "Transparent Initiator Mode" section on page 45-16), if an iSCSI host has static WWN mapping then the standard Fibre Channel device pWWN-based zoning membership mechanism can be used.

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

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

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

Note In proxy initiator mode, all iSCSI devices connecting to an IPS port gain access to the Fibre Channel fabric through a single virtual Fibre Channel N port. Thus, zoning based on the iSCSI node name or IP address will not have any effect. If zoning based on pWWN is used, then all iSCSI devices connecting to that IPS port will be put in the same zone. To implement individual initiator access control in proxy initiator mode, configure an iSCSI ACL on the virtual target (see the "iSCSI-Based Access Control" section on page 45-29).

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

Step 1 Select **Zone > Edit Local Full Zone Database**.

You see the Edit Local Zone Database dialog box.

Figure 45-32 Edit Local Zone Database Dialog Box in Fabric Manager

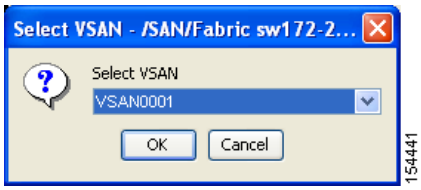

Step 2 Select the VSAN you want to add the iSCSI host initiator to and click **OK**.

You see the available zones and zone sets for that VSAN (see Figure 45-33).

Figure 45-33 Available Zones and Zone Sets

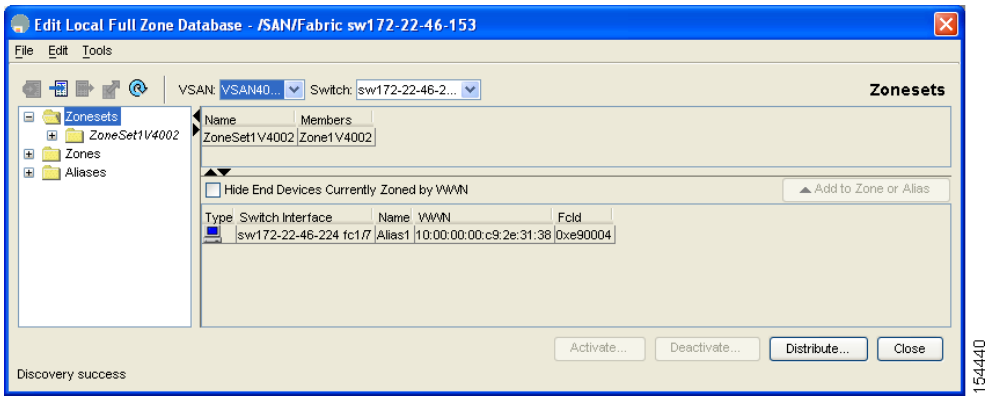

Step 3 From the list of available devices with iSCSI host initiators, drag the initiators to add into the zone.

Step 4 Click **Distribute** to distribute the change.

Step 5 Click **Close** to close the dialog box.

iSCSI-Based Access Control

iSCSI-based access control is applicable only if static iSCSI virtual targets are created (see the "Static Mapping" section on page 45-10). For a static iSCSI target, you can configure a list of iSCSI initiators that are allowed to access the targets.

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

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

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

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

Step 1 Select **IP > iSCSI.**

You see the iSCSI configuration dialog box.

Figure 45-34 iSCSI Tables in Device Manager

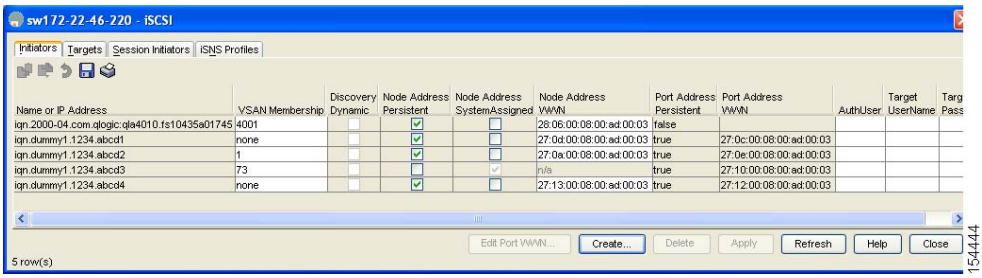

Step 2 Select the **Targets** tab.

You see the iSCSI virtual targets.

Step 3 Uncheck the **Initiators Access All** check box if checked.

Step 4 Click **Edit Access**.

You see the Initiators Access dialog box.

Step 5 Click **Create** to add more initiators to the Initiator Access list.

You see the Create Initiators Access dialog box.

Step 6 Add the name or IP address for the initiator that you want to permit for this virtual target.

Step 7 Click **Create** to add this initiator to the Initiator Access List or click **Close** to discard any unsaved changes.

Enforcing Access Control

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

- iSCSI discovery phase—When an iSCSI host creates an iSCSI discovery session and queries for all iSCSI targets, the IPS module or MPS-14/2 module returns only the list of iSCSI targets this iSCSI host is allowed to access based on the access control policies discussed in the previous section. The IPS module or MPS-14/2 module does this by querying the Fibre Channel name server for all the devices in the same zone as the initiator in all VSANs. It then filters out the devices that are initiators by looking at the FC4-feature field of the FCNS entry. (If a device does not register as either initiator or target in the FC4-feature field, the IPS module or MPS-14/2 module will advertise it.) It then responds to the iSCSI host with the list of targets. Each will have either a static iSCSI target name that you configure or a dynamic iSCSI target name that the IPS module or MPS-14/2 module creates for it (see the "Dynamic Mapping" section on page 45-8).
- iSCSI session creation—When an IP host initiates an iSCSI session, the IPS module or MPS-14/2 module verifies if the specified iSCSI target (in the session login request) is allowed by both the access control mechanisms described in the "iSCSI-Based Access Control" section on page 45-29.

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

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

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

iSCSI Session Authentication

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

For CHAP user name or secret validation, you can use any method supported and allowed by the Cisco MDS AAA infrastructure (see Chapter 35, "Configuring RADIUS and TACACS+"). AAA authentication supports a RADIUS, TACACS+, or local authentication device.

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

Step 1 Choose **Switches > Security > AAA** from the Physical Attributes pane.

You see the AAA configuration in the Information pane.

Step 2 Select the **Applications** tab.

You see the AAA configuration per application.

Figure 45-35 AAA per Application Configuration

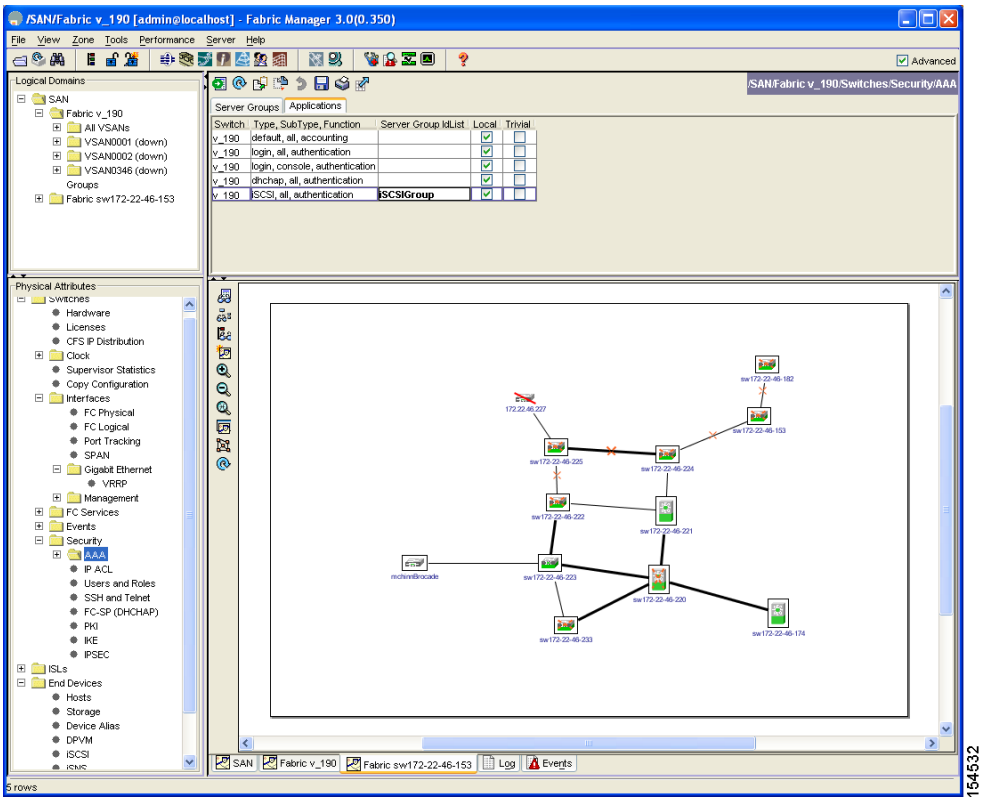

Step 3 Right-click on the ServerGroup Id List field for the iSCSI application and enter the server group that you want iSCSI to use.

Note You should use an existing server group or create a new server group before configuring it for iSCSI session authentication.

Step 4 Click the **Apply Changes icon** to save these changes or click **Undo Changes** to discard any unsaved changes.

Authentication Mechanism

You can configure iSCSI CHAP or None authentication at both the global level and at each interface level.

The authentication for a Gigabit Ethernet interface or subinterface overrides the authentication method configured at the global level.

To configure the authentication mechanism for iSCSI using Fabric Manager, follow these steps:

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

You see the iSCSI tables in the Information pane.

Figure 45-36 iSCSI Tables in Fabric Manager

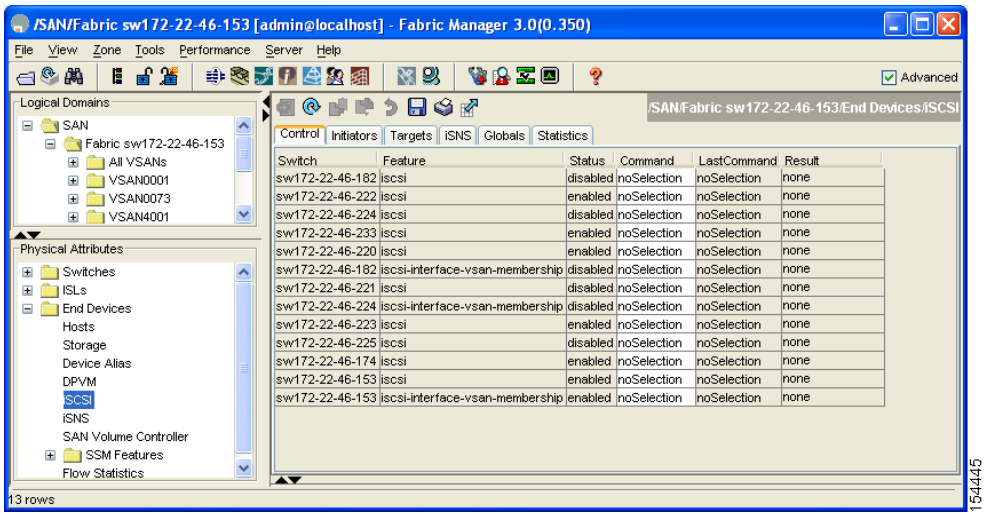

Step 2 Click the **Global** tab.

You see the iSCSI authentication configuration table.

Step 3 Select **chap** or **none** from the authMethod column.

Step 4 Click the **Apply Changes** icon in Fabric Manager to save these changes, or click the **Undo Changes** icon to discard changes.

To configure the authentication mechanism for iSCSI sessions to a particular interface using Fabric Manager, follow these steps:

Step 1 Choose **Switches > Interfaces > Gigabit Ethernet** from the Physical Attributes pane. You see the Gigabit Ethernet configuration in the Information pane.

Step 2 Select the **iSNS** tab. You see the iSCSI and iSNS configuration.

Figure 45-37 Configuring iSCSI Authentication on an Interface

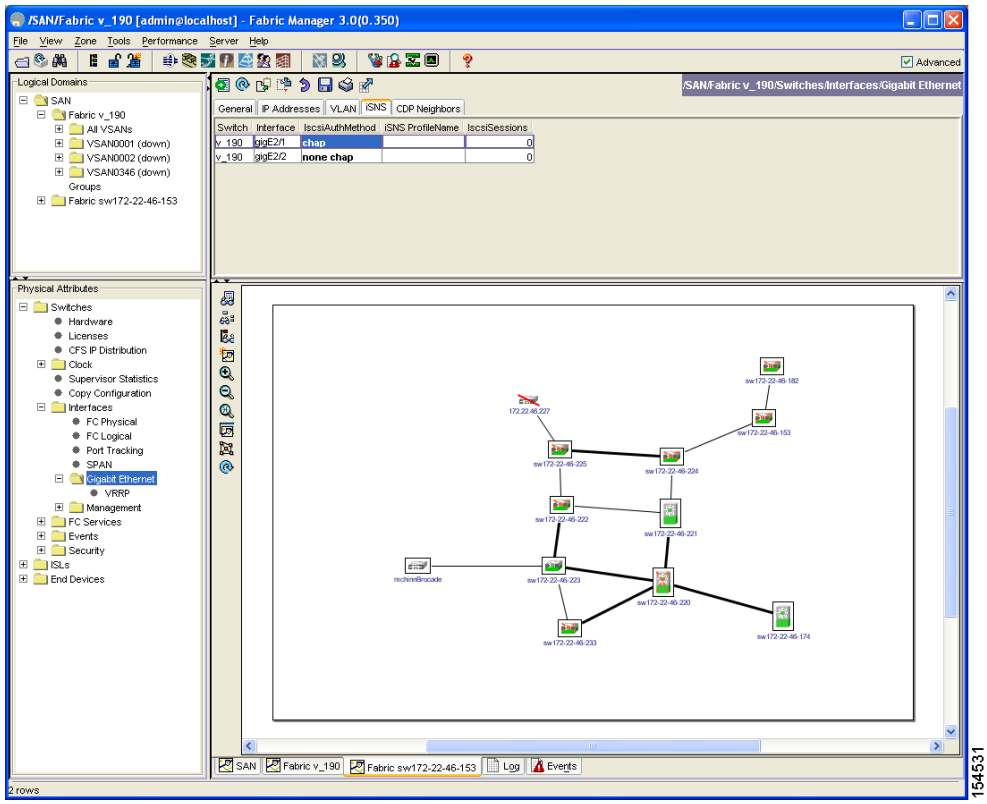

Step 3 Right-click on the **IscsiAuthMethod** field and select none or chap.

Step 4 Click the **Apply Changes** icon to save these changes, or click the **Undo Changes** icon to discard changes.

Local Authentication

See the "Configuring Users" section on page 33-12 to create the local password database. To create users in the local password database for the iSCSI initiator, the iSCSI keyword is mandatory.

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

Step 1 Choose **Security > iSCSI**.

You see the iSCSI Security dialog box. See Figure 45-38.

Figure 45-38 Configuring iSCSI Users for Local Authentication

Step 2 Complete the iSCSI User, Password, and Password Confirmation fields.

Step 3 Click **Create** to save this new user or click **Close** to discard any unsaved changes.

Restricting iSCSI Initiator Authentication

By default, the iSCSI initiator can use any user name in the RADIUS server or in the local database in authenticating itself to the IPS module or MPS-14/2 module (the CHAP user name is independent of the iSCSI initiator name). The IPS module or MPS-14/2 module allows the initiator to login as long as it provides a correct response to the CHAP challenge sent by the switch. This can be a problem if one CHAP user name and password had been compromised.

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

Step 1 Exapnd **End Devices** and then select **iSCSI** from the Physical Attributes pane.

You see the iSCSI tables in the Information pane.

Figure 45-39 iSCSI Tables in Fabric Manager

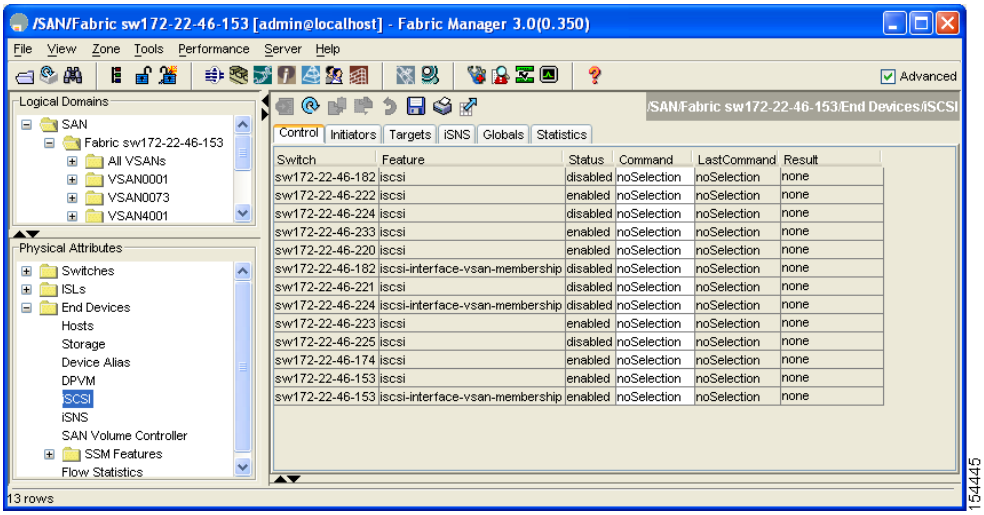

Step 2 Right-click the AuthUser field and enter the user name that you want to restrict the iSCSI initiator to.

Step 3 Click the **Apply Changes** icon to save these changes or click the **Undo Changes icon** to discard any unsaved changes.

Mutual CHAP Authentication

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

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

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

You see the iSCSI tables in the Information pane.

Figure 45-40 iSCSI Tables in Fabric Manager

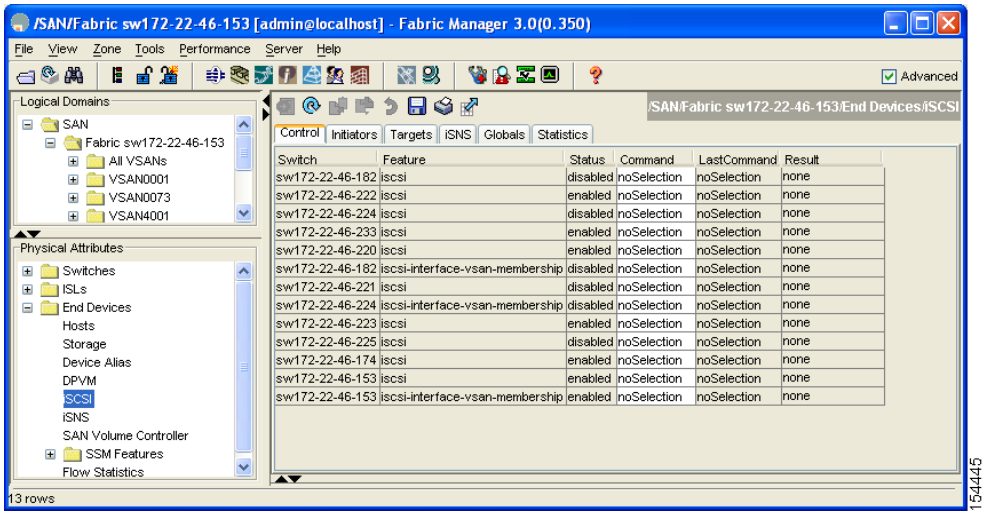

Step 2 Select the **Globals** tab.

You see the global iSCSI configuration.

Step 3 Fill in the Target UserName and Target Password fields.

Step 4 Click the **Apply Changes** icon to save these changes or click the **Undo Changes icon** to discard any unsaved changes.

To configure a per-initiator iSCSI target's user name and password used by the switch to authenticate itself to an initiator using Device Manager, follow these steps:

Step 1 Select **IP > iSCSI.**

You see the iSCSI configuration dialog box.

Figure 45-41 iSCSI Tables in Device Manager

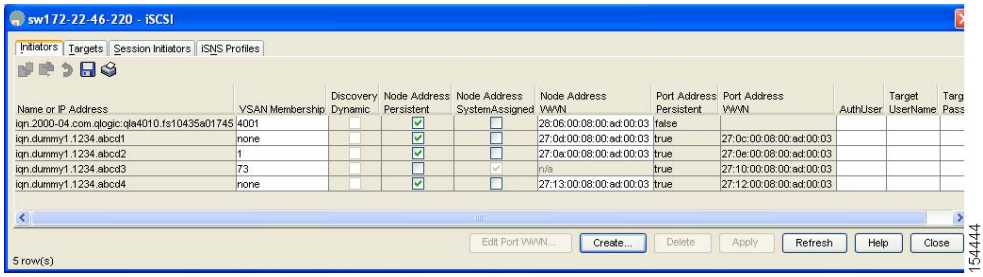

Step 2 Complete the Target UserName and Target Password fields for the initiator that you want to configure.

Step 3 Click **Create** to add this initiator to the Initiator Access List or click **Close** to discard any unsaved changes.

Configuring an iSCSI RADIUS Server

To configure an iSCSI RADIUS server, follow these steps:

Step 1 Configure the RADIUS server to allow access from the Cisco MDS switch's management Ethernet IP address.

Step 2 Configure the shared secret for the RADIUS server to authenticate the Cisco MDS switch.

Step 3 Configure the iSCSI users and passwords on the RADIUS server.

iSCSI Immediate Data and Unsolicited Data Features

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

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

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

iSCSI Interface Advanced Features

Advanced configuration options are available for iSCSI interfaces on a per-IPS port basis. These configurations are similar to the advanced FCIP configurations and are already explained in that section (see the "Advanced FCIP Interface Configuration" section on page 43-17).

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

iSCSI Listener Port

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

See the "Configuring TCP Listener Ports" section on page 43-14.

TCP Tuning parameters

You can configure the following TCP parameters:

- Minimum retransmit timeout (See the "Minimum Retransmit Timeout" section on page 43-14.)
- Keepalive timeout (See the "Keepalive Timeout" section on page 43-14.)
- Maximum retransmissions (See the "Maximum Retransmissions" section on page 43-15)
- Path MTU (See the "Path MTUs" section on page 43-15.)
- SACK (SACK is enabled by default for iSCSI TCP configurations.) (See the "Selective Acknowledgments" section on page 43-15.)
- Window management (The iSCSI defaults are max-bandwidth is 1 Gbps, min-available-bandwidth is 70 Mbps, and round-trip-time is 1 msec.) (See the "Window Management" section on page 43-15.)
- Buffer size (The iSCSI default send buffer size is 4096 KB) (See the "Buffer Size" section on page 43-16.)
- Window congestion monitoring (enabled by default and the default burst size is 50 KB) (See the "Monitoring Congestion" section on page 43-16.)

  Maximum delay jitter (enabled by default and the default time is 500 microseconds) (See the "Estimating Maximum Jitter" section on page 43-16.)

QoS Values

To set the QoS values, follow these steps:

Step 1 In Fabric Manager, expand **Switches**, expand **Interfaces** and then select **FC Logical** from the Physical Attributes pane.

You see the Interface tables in the Information pane.

Figure 45-42 FC Logical Interface Tables

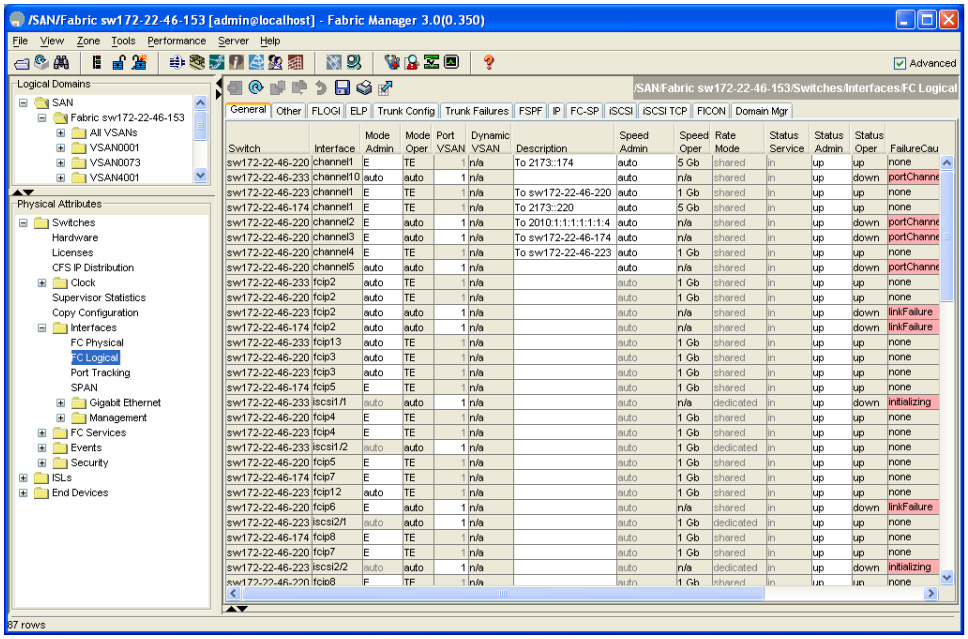

In Device Manager, click **Interfaces > Ethernet and iSCSI**.

You see the interfaces dialog box.

Figure 45-43 Ethernet and iSCSI Interfaces in Device Manager

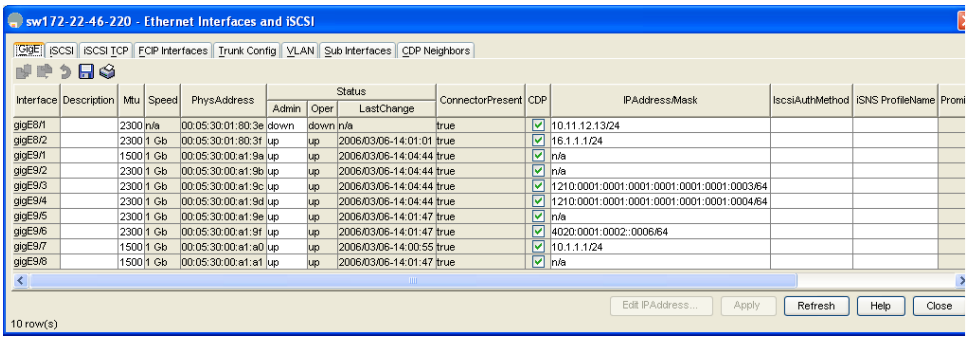

Step 2 Click the **iSCSI TCP** tab in either Fabric Manager or Device Manager.

You see the iSCSI TCP configuration table.

Step 3 Set the QoS field from 1 to 6.

Step 4 Click the **Apply Changes** icon in Fabric Manager or click **Apply** in Device Manager to save these changes, or click **Undo Changes** in Fabric Manager or click **Cancel** in Device Manager to discard changes.

iSCSI Routing Modes

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

- Pass-thru mode

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

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

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

- Store-and-forward mode (default)

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

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

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

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

- Cut-through mode

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

Figure 45-44 compares the messages exchanged by the iSCSI routing modes.

Figure 45-44 iSCSI Routing Modes

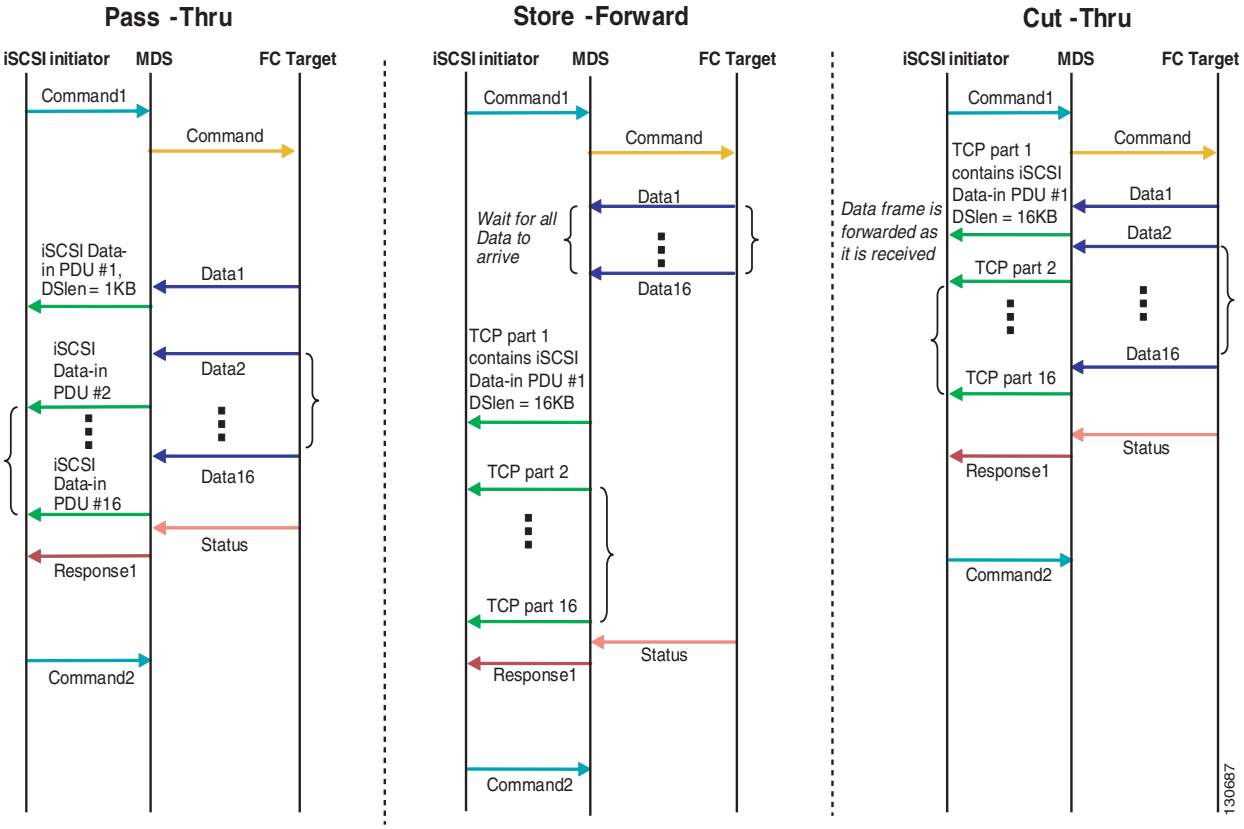

Table 45-1 compares the advantages and disadvantages of the different iSCSI routing modes.

Table 45-1 Comparison of iSCSI Routing Modes

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

Caution Changing the forwarding mode of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the "Changing iSCSI Interface Parameters and the Impact on Load Balancing" section on page 45-52.

About iSLB

The iSCSI server load balancing (iSLB) feature provides a means to easily configure large scale iSCSI deployments containing hundreds or even thousands of initiators. When not using iSLB, configuring iSCSI requires the following:

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

iSLB provides the following features:

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

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

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

Configuring iSLB

This section covers the following topics:

- Distributing the iSLB Configuration Using CFS, page 45-53
- iSCSI High Availability, page 45-58

Note Before configuring iSLB, you must enable iSCSI (see the "Enabling iSCSI" section on page 45-4).

About iSLB Configuration Limits

iSLB configuration has the following limits:

- The maximum number of iSLB initiators supported in a fabric is 2000.
- The maximum number of iSLB sessions per IPS port in either transparent or proxy initiator mode is 500.
- The maximum number of switches in a fabric that can have iSLB with CFS distribution enabled is four.
- No more than 200 new iSLB initiators can be added to the pending configuration. Before adding more initiators, you must commit the configuration.
- You cannot disable iSCSI if you have more than 200 iSLB initiators in the running configuration. Reduce the number of iSLB initiators to fewer than 200 before disabling iSCSI.
- iSLB can be used without CFS distribution but if iSLB auto-zone feature is used, traffic is disrupted when any zoneset is activated.
- If IVR and iSLB features are enabled in the same fabric, you should have at least one switch in the fabric where both these features are enabled. Any zoning-related configuration and activation (for normal zones, IVR zones, or iSLB zones) must be performed on this switch. Otherwise, there may be traffic disruption in the fabric.

iSLB Configuration Prerequisites

Do the following prior to configuring iSLB:

- Enable iSCSI (see the "Enabling iSCSI" section on page 45-4).
- Configure the Gigabit Ethernet interfaces (see the "Configuring Gigabit Ethernet Interfaces for IPv4" section on page 47-4 or the Configuring Basic Connectivity for IPv6, page 48-21).
- Configure the VRRP groups (see the "Configuring Load Balancing Using VRRP" section on page 45-52).
- Configure and activate a zone set (see Chapter 26, "Configuring and Managing Zones").
- Enable CFS distribution for iSLB.

About iSLB Initiators

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

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

Configuring iSLB Initiators

This section includes the following topics:

Assigning WWNs to iSLB Initiators

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

- Dynamic mapping (default)
- Static mapping

Note Assigning WWNs for iSLB initiators is the same as for iSCSI initiators. For information on dynamic and static mapping, see the "WWN Assignment for iSCSI Initiators" section on page 45-17.

Tip We recommend using the **SystemAssign** option. If you manually assign a WWN, you must ensure its uniqueness (see the "Configuring World Wide Names" section on page 62-22). You should not use any previously-assigned WWNs.

See the "Configuring iSLB using Device Manager" procedure on page 45-47.

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

Making the Dynamic iSLB Initiator WWN Mapping Static

After a dynamic iSLB initiator has logged in, you may decide to permanently keep the automatically assigned nWWN/pWWN mapping to allow this initiator to use the same mapping the next time it logs in (see the "Dynamic Mapping" section on page 45-8).

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

Note You cannot convert a dynamic iSCSI initiator to a static iSLB initiator or a dynamic iSLB initiator to a static iSCSI initiator.

Note Making the dynamic mapping for iSLB initiators static is the same as for iSCSI. See the "Making the Dynamic iSLB Initiator WWN Mapping Static" section on page 45-45.

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

See the "Configuring iSLB using Device Manager" procedure on page 45-47.

Assigning VSAN Membership for iSLB Initiators

Individual iSLB hosts can be configured to be in a specific VSAN (similar to the DPVM feature for Fibre Channel; see Chapter 2, "Installing Cisco MDS SAN-OS and Fabric Manager"). The specified VSAN overrides the iSCSI interface VSAN membership.

Note Specifying the iSLB initiator VSAN is the same as for an iSCSI initiator. See the "VSAN Membership for iSCSI" section on page 45-24.

Note When an iSLB initiator is configured in any other VSAN (other than VSAN 1, the default VSAN), for example VSAN 2, the initiator is automatically removed from VSAN 1. If you also want it to be present in VSAN 1, you must explicitly configure the initiator in VSAN 1.

See the "Configuring iSLB using Device Manager" procedure on page 45-47.

Configuring Metric for Load Balancing

You can assign a load metric to each initiator for weighted load balancing. The load calculated is based on the number of initiators on a given iSCSI interface. This feature accommodates initiators with different bandwidth requirements. For example, you could assign a higher load metric to a a database server than to a web server. Weighted load balancing also accommodates initiators with different link speeds.

For more information on load balancing, see the "Configuring iSLB Initiator Targets" section on page 45-49.

Click **IP > iSCSI iSLB** in Device Manager and set the LoadMetric field to change the load balancing metric for an iSLB initiator.

See the "Configuring iSLB using Device Manager" procedure on page 45-47.

Configuring and Activating Zones for iSLB Initiators and Initiator Targets

You can configure a zone name where the iSLB initiators and initiator targets are added. If you do not specify a zone name, the IPS manager creates one dynamically. iSLB zone sets have these considerations:

- Auto-zoning of the initiator with the initiator targets is enabled by default.
- A zone set must be active in a VSAN for auto-zones to be created in that VSAN.
- iSLB zone set activation might fail if another zone set activation is in process or if the zoning database is locked. Retry the iSLB zone set activation if a failure occurs. To avoid this problem, only perform only one zoning related operation (normal zones, IVR zones, or iSLB zones) at a time.
- If IVR and iSLB are enabled in the same fabric, at least one switch in the fabric must have both features enabled. Any zoning related configuration or activation operation (for normal zones, IVR zones, or iSLB zones) must be performed on this switch. Otherwise, traffic might be disrupted in the fabric.

Choose **IP > iSCSI iSLB** in Device Manager and set the autoZoneName field to change the auto zone name for an iSLB initiator.

See the "Configuring iSLB using Device Manager" procedure on page 45-47.

Configuring iSLB Session Authentication

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

For CHAP user name or secret validation you can use any method supported and allowed by the Cisco MDS AAA infrastructure (see Chapter 35, "Configuring RADIUS and TACACS+"). AAA authentication supports RADIUS, TACACS+, or a local authentication device.

Note Specifying the iSLB session authentication is the same as for iSCSI. See the "iSCSI Session Authentication" section on page 45-31.

Restricting iSLB Initiator Authentication

By default, the iSLB initiator can use any user name in the RADIUS or local AAA database in authenticating itself to the IPS module or MPS-14/2 module (the CHAP user name is independent of the iSLB initiator name). The IPS module or MPS-14/2 module allows the initiator to log in as long as it provides a correct response to the CHAP challenge sent by the switch. This can be a problem if one CHAP user name and password have been compromised.

Choose **IP > iSCSI iSLB** in Device Manager and set the AuthName field to restrict an initiator to use a specific user name for CHAP authentication.

See the "Configuring iSLB using Device Manager" procedure on page 45-47.

Mutual CHAP Authentication

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

Choose **IP > iSCSI iSLB** in Device Manager and set the Target Username and Target Password fields to configure a per-initiator user name and password used by the switch to authenticate itself to an initiator.

See the "Configuring iSLB using Device Manager" procedure on page 45-47.

Configuring iSLB using Device Manager

To configure iSLB using Device Manager, follow these steps:

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

You see the iSCSI iSLB Configuration dialog box.

Figure 45-45 iSLB Configuration Dialog Box in Device Manager

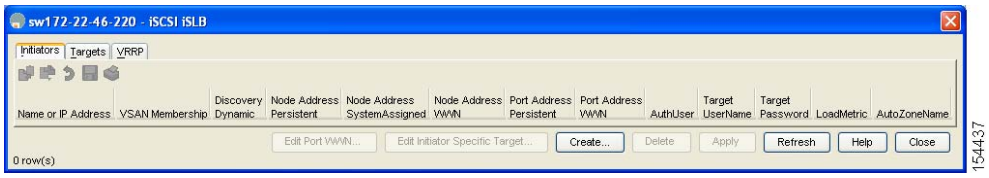

Step 2 Click **Create** to create a new iSCSI iSLB initiator.

You see the Create iSCSI iSLB dialog box shown in Figure 45-46.

Figure 45-46 Create iSLB Dialog Box in Device Manager

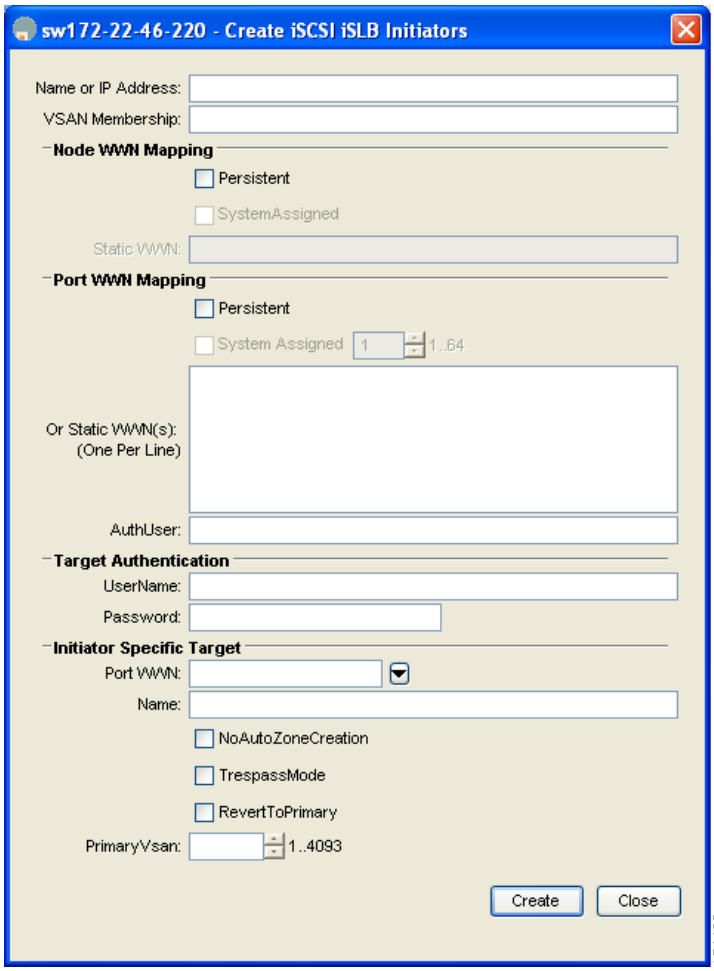

Step 3 Set the Name or IP Address field to the iSLB name or IP address.

Step 4 Set the VSAN Membership field to the VSAN that you want the iSLB initiator in (see Figure 45-46). Also see the "Assigning VSAN Membership for iSLB Initiators" section on page 45-45.

Step 5 Check the **Persistent** check box to convert a dynamic nWWN to static for the iSLB initiator (see Figure 45-46). Also see the "Making the Dynamic iSLB Initiator WWN Mapping Static" section on page 45-45.

Step 6 Optionally, check the **SystemAssigned** check box to have the switch assign the nWWN (see Figure 45-46). Also see the "Assigning WWNs to iSLB Initiators" section on page 45-44.

Step 7 Optionally, set the Static WWN field to manually assign the static nWWN. You must ensure uniqueness for this nWWN.

Step 8 Optionally, check the Port WWN Mapping **Persistent** check box convert dynamic pWWNs to static for the iSLB initiator. See the "Making the Dynamic iSLB Initiator WWN Mapping Static" section on page 45-45.

Step 9 Optionally, check the **SystemAssigned** check box and set the number of pWWNs you want to have the switch assign the PWWN (see Figure 45-46). Also see the "Assigning WWNs to iSLB Initiators" section on page 45-44.

Step 10 Optionally, set the Static WWN(s) field to manually assign the static pWWNs (see Figure 45-46). You must ensure uniqueness for these pWWN.

Step 11 Optionally, set the AuthUser field to the username that you want to restrict the iSLB initiator to for iSLB authentication (see Figure 45-46). Also see the "Restricting iSLB Initiator Authentication" section on page 45-47.

Step 12 Fill in the Username and Password fields to configure iSLB initiator target CHAP authentication (see Figure 45-46). Also see the "Configuring iSLB Session Authentication" section on page 45-46.

Step 13 In the Initiator Specific Target section, set the pWWN to configure an iSLB initiator target (see Figure 45-46). Also see the "Configuring iSLB Initiator Targets" section on page 45-49.

Step 14 Optionally, set the Name field to a globally unique identifier (IQN).

Step 15 Optionally, check the **NoAutoZoneCreation** check box to disable auto-zoning (see Figure 45-46). Also see the "Configuring and Activating Zones for iSLB Initiators and Initiator Targets" section on page 45-46.

Step 16 Optionally, check the **TresspassMode** check box (see Figure 45-46). Also see the "LUN Trespass for Storage Port Failover" section on page 45-61.

Step 17 Optionally, check the **RevertToPrimary** check box to revert back to the primary port after an HA failover when the primary port comes back up (see Figure 45-46).

Step 18 Set the PrimaryVsan to the VSAN for the iSLB initiator target.

Step 19 Click **Create** to create this iSLB initiator or click Close to close the dialog box without saving any changes.

Step 20 If CFS is enabled, select **commit** from the CFS drop-down menu.

Configuring iSLB Initiator Targets

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

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

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

In addition, you can disable auto-zoning.

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

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

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

You see the iSCSI iSLB Configuration dialog box.

Figure 45-47 iSLB Configuration Dialog Box in Device Manager

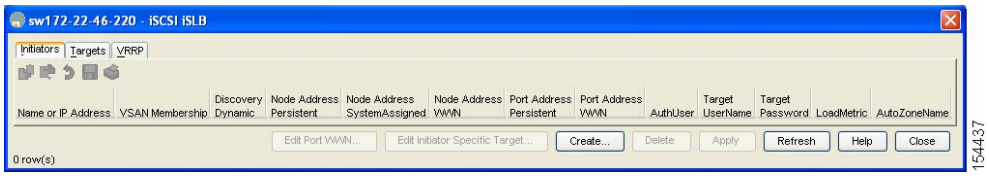

Step 2 Right-click on the initiator you want to add targets to and click **Edit Initiator Specific Targets**.You see the Initiator Specific Targets Configuration dialog box.

Step 3 Click **Create** to create a new initiator target. You see the Create Initiator Specific Target dialog box.

Step 4 Fill in the pWWN field with the initiator target pWWN.

Step 5 Optionally, set the Name field to a globally unique identifier (IQN).

Step 6 Optionally, check the **NoAutoZoneCreation** check box to disable auto-zoning (see Figure 45-46). Also see the "Configuring and Activating Zones for iSLB Initiators and Initiator Targets" section on page 45-46.

Step 7 Optionally, check the **TresspassMode** check box. See the "LUN Trespass for Storage Port Failover" section on page 45-61.

Step 8 Optionally, check the **RevertToPrimary** check box to revert back to the primary port after an HA failover when the primary port comes back up.

Step 9 Set the PrimaryVsan to the VSAN for the iSLB initiator target.

Step 10 Click **Create** to create this iSLB initiator target or click Close to close the dialog box without saving any changes.

Step 11 If CFS is enabled, select **commit** from the CFS drop-down menu.

About Load Balancing Using VRRP

You can configure Virtual Router Redundancy Protocol (VRRP) load balancing for iSLB. Figure 45-48 shows an example of load balancing using iSLB.

Figure 45-48 iSLB Initiator Load Balancing Example

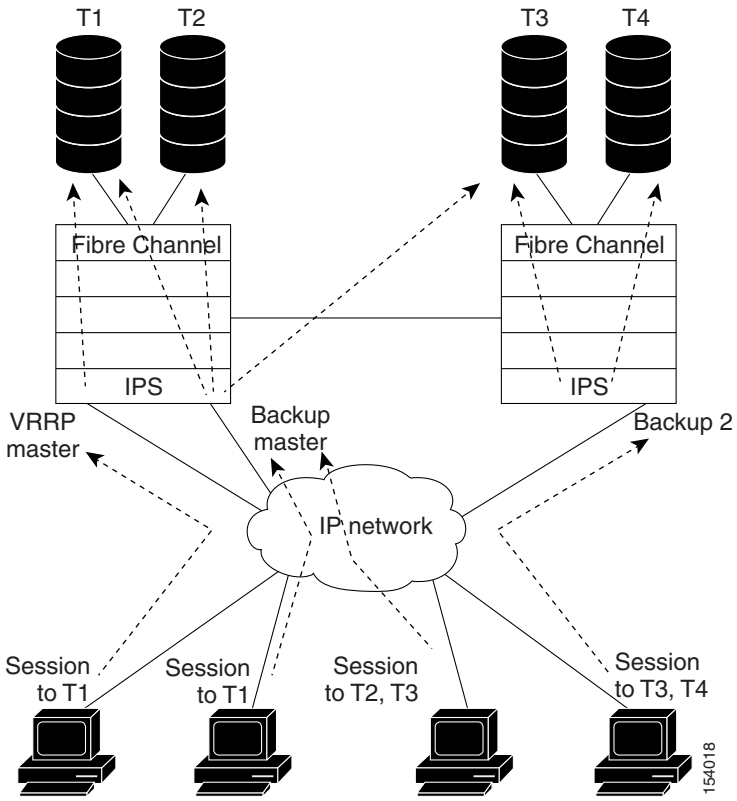

The host is configured with a VRRP address as the portal address. When the VRRP master port receives the first iSCSI session from an initiator, it assigns a backup port to serve that particular host. The information is synchronized to all switches through CFS if recovery is needed when a master port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

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

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

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

Changing iSCSI Interface Parameters and the Impact on Load Balancing

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

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

VRRP Load Balancing Algorithm For Selecting Gigabit Ethernet Interfaces

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

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

VRRP backup interface load > [2 * VRRP master interface load + 1]

Configuring Load Balancing Using VRRP

You must first configure VRRP on the Gigabit Ethernet interfaces on the switch that connect to the IP network before configuring VRRP for iSLB. For information on how to configure VRRP on a Gigabit Ethernet interface, see the "Virtual Router Redundancy Protocol" section on page 46-8.

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

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

You see the iSLB configuration dialog box.

Figure 45-49 iSLB Configuration Dialog Box in Device Manager

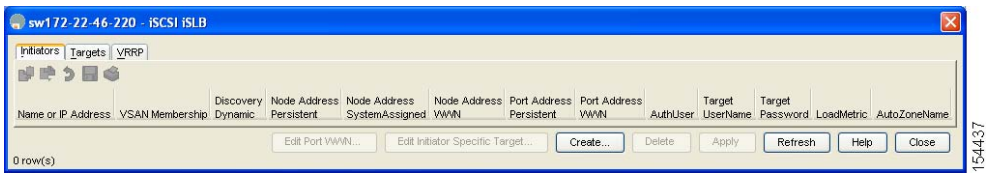

Step 2 Select the **VRRP** tab and click **Create** to configure VRRP load balancing for iSLB initiators.

You see the Create VRRP Dialog Box shown in Figure 45-50.

Figure 45-50 Create VRRP Dialog Box

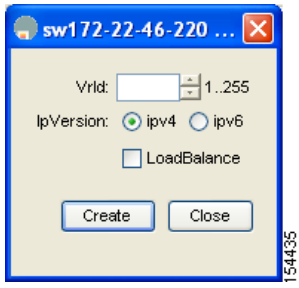

Step 3 Set the Vrld to the VRRP group number.

Step 4 Select either **ipv4** or **ipv6** and check the LoadBalance check box.

Step 5 Click **Create** to enable load balancing or click **Close** to close the dialog box without saving any changes.

Step 6 If CFS is enabled, select **commit** from the CFS drop-down menu.

About iSLB Configuration Distribution Using CFS

Configuration for iSLB initiators and initiator targets on an MDS switch can be distributed using the Cisco Fabric Services (CFS). This feature allows you to synchronize the iSLB configuration across the fabric from the console of a single MDS switch. The iSCSI initiator idle timeout, global authentication, and iSCSI dynamic initiator mode parameters are also distributed. CFS distribution is disabled by default (see Chapter 5, "Using the CFS Infrastructure").

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

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

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

Non-iSLB Virtual Targets will continue to support advertised interfaces option.

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

Distributing the iSLB Configuration Using CFS

This section contains the following:

Enabling iSLB Configuration Distribution

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

Step 1 Choose **Admin > CFS**. You see the CFS Configuration dialog box (see Figure 45-51).

Figure 45-51 Enabling CFS in Device Manager

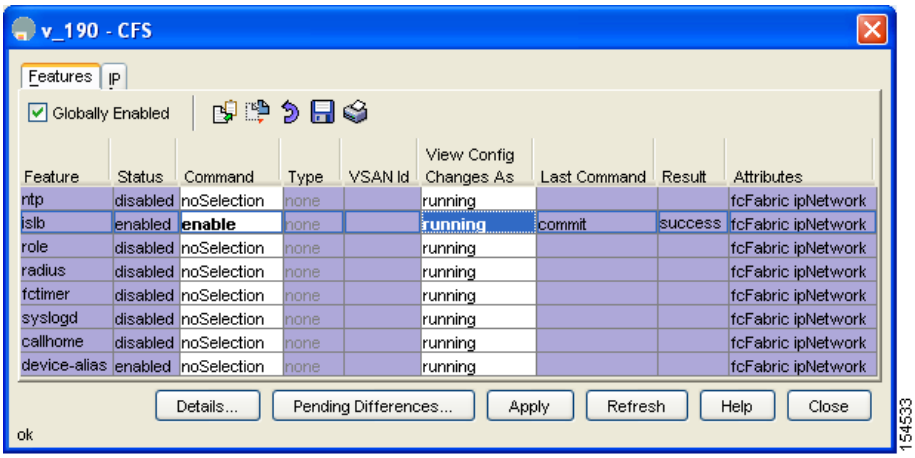

Step 2 Set the Command field to **enable** for the iSLB feature.

Step 3 Click **Apply** to save this change or click **Close** to discard any unsaved changes.

Locking the Fabric

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

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

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

Committing Changes to the Fabric

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

To commit changes to the fabric using Device Manager, follow these steps:

Step 1 Choose **Admin > CFS**. You see the CFS Configuration dialog box (see Figure 45-51).

Figure 45-52 Enabling CFS in Device Manager

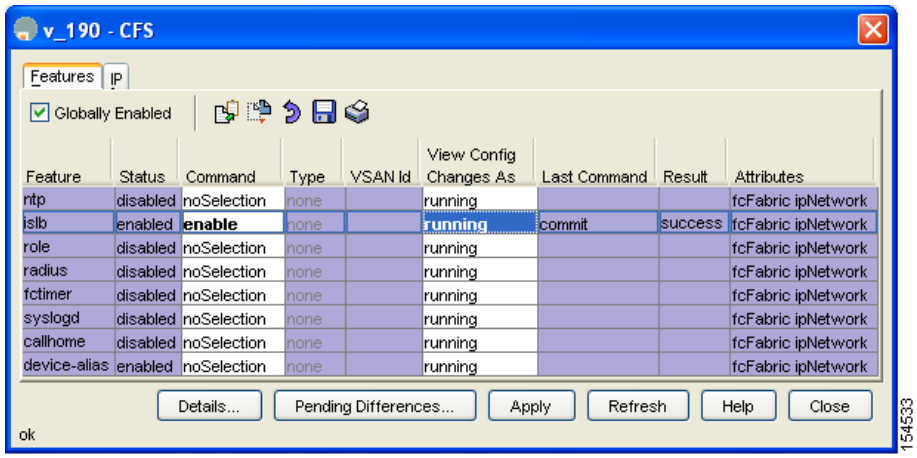

Step 2 Set the Command field to **commit** for the iSLB feature.

Step 3 Click **Apply** to save this change or click **Close** to discard any unsaved changes.

Discarding Pending Changes

At any time, you can discard the pending changes to the iSLB configuration and release the fabric lock. This action has no affect on the active configuration on any switch in the fabric.

To discard unsaved changes to the fabric using Device Manager, follow these steps:

Step 1 Choose **Admin > CFS**.

You see the CFS Configuration dialog box (see Figure 45-51).

Figure 45-53 Enabling CFS in Device Manager

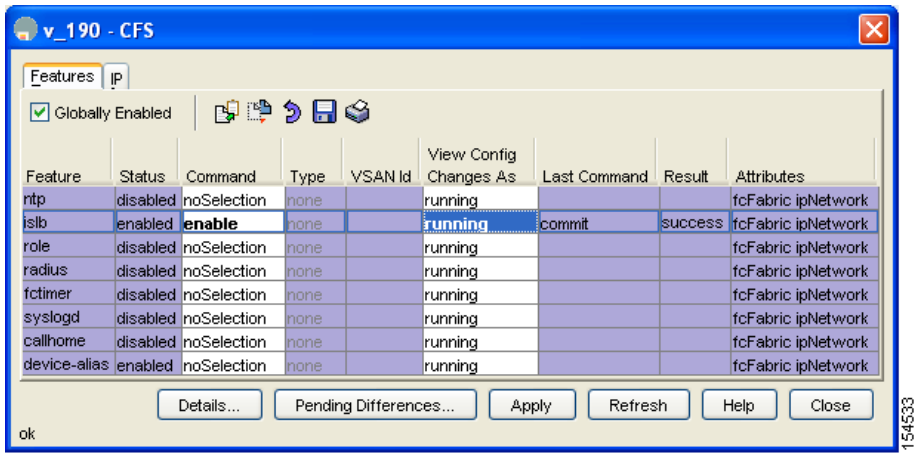

Step 2 Set the Command field to **abort** for the iSLB feature.

Step 3 Click **Apply** to save this change or click **Close** to discard any unsaved changes.

Clearing a Fabric Lock

If you have performed an iSLB configuration task and have not released the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your pending changes are discarded and the fabric lock is released.

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

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

Step 1 Choose **Admin > CFS**. You see the CFS Configuration dialog box (see Figure 45-51).

Figure 45-54 Enabling CFS in Device Manager

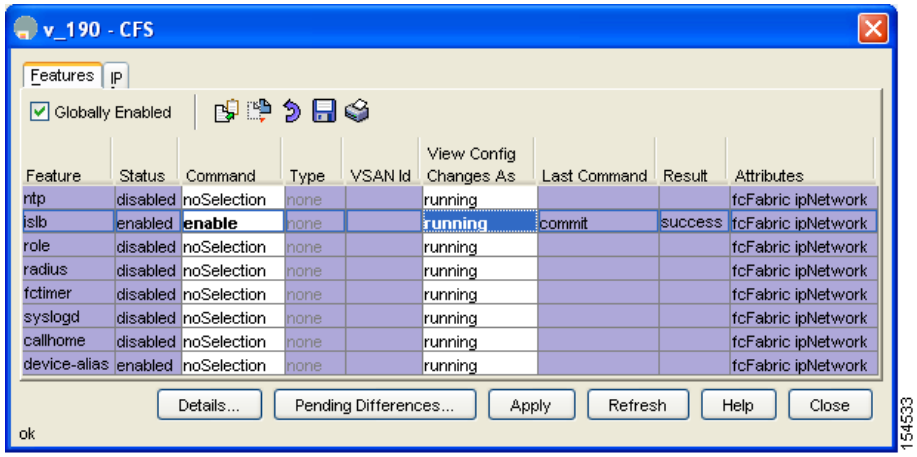

Step 2 Set the Command field to **clear** for the iSLB feature.

Step 3 Click **Apply** to save this change or click **Close** to discard any unsaved changes.

iSLB CFS Merge Status Conflicts

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

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

Tip Check the syslog for details on merge conflicts.

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

- Same Virtual Target name with different configuration on different switches

- Different Virtual Target name with same target pwwn configuration on different switches

  Allowed today on a single switch

  Not entirely desired if the merge creates such targets as it may let the initiators see the same FC target multiple times

  Handled by LUN overlap check, so only one session will come up

iSCSI High Availability

The following high availability features are available for iSCSI configurations:

- Transparent Target Failover, page 45-58
- Multiple IPS Ports Connected to the Same IP Network, page 45-62
- VRRP-Based High Availability, page 45-63
- Ethernet PortChannel-Based High Availability, page 45-64

Transparent Target Failover

The following high availability configurations are available:

- iSCSI high availability with host running multi-path software
- iSCSI High availability with host not having multi-path software

iSCSI High Availability with Host Running Multi-Path Software

Figure 45-55 shows the physical and logical topology for an iSCSI HA solution for hosts running multi-path software. In this scenario, the host has four iSCSI sessions. There are two iSCSI sessions from each host NIC to the two IPS ports.

Figure 45-55 Host Running Multi-Path Software

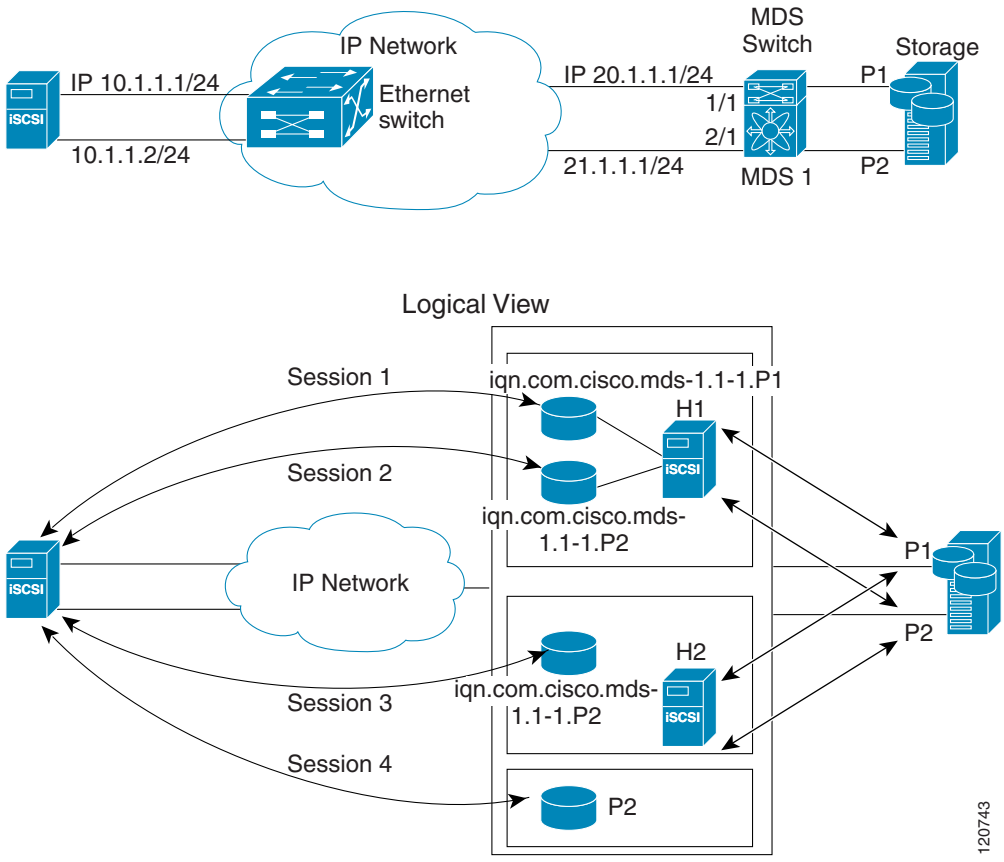

Each IPS ports is exporting the same two Fibre Channel target ports of the storage but as different iSCSI target names if you use dynamic iSCSI targets). So the two IPS ports are exporting a total of four iSCSI target devices. These four iSCSI targets map the same two ports of the Fibre Channel target.

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

If the iSCSI host NIC-1 fails (see Figure 45-55 for the physical view), then sessions 1 and 2 fail but we still have sessions 3 and 4.

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

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

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

iSCSI HA with Host Not Having Any Multi-Path Software

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

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

- IPS ports support the VRRP feature (see the "Configuring VRRP for Gigabit Ethernet Interfaces" section on page 47-9) to provide failover for IPS ports.
- IPS has transparent Fibre Channel target failover for iSCSI static virtual targets.

Statically imported iSCSI targets have an additional option to provide a secondary pWWN for the Fibre Channel target. This can be used when the physical Fibre Channel target is configured to have an LU visible across redundant ports. When the active port fails, the secondary port becomes active and the iSCSI session switches to use the new active port (see Figure 45-56).

Figure 45-56 Static Target Importing Through Two Fibre Channel Ports

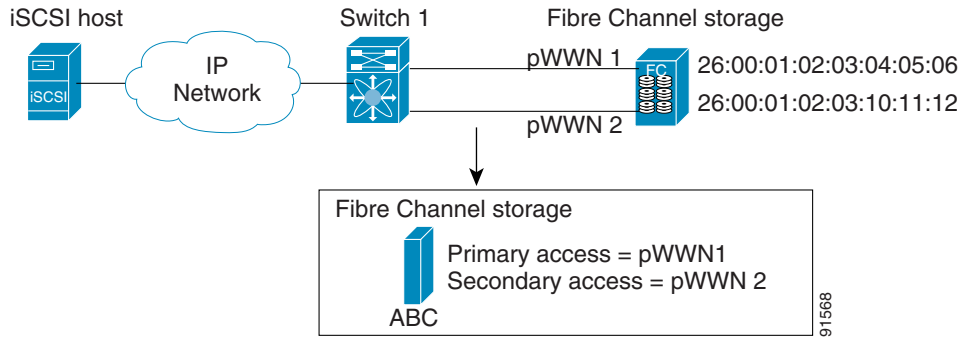

In Figure 45-56, you can create an iSCSI virtual target that is mapped to both pWWN1 and pWWN2 to provide redundant access to the Fibre Channel targets.

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

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

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

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

Step 1 Click **IP > iSCSI**.

You see the iSCSI tables under the Initiator tab.

Step 2 Click the **Targets** tab to display a list of existing iSCSI targets shown in Figure 45-57.

Figure 45-57 iSCSI Targets in Device Manager

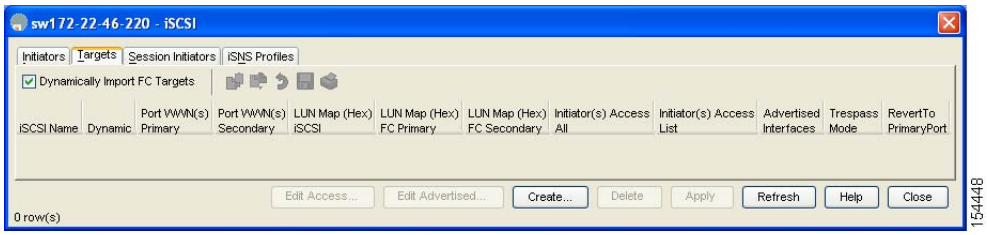

Step 3 Click **Create** to create an iSCSI target. You see the Create iSCSI Targets dialog box.

Step 4 Set the iSCSI target node name in the iSCSI Name field, in IQN format.

Step 5 Set the Port WWN field for the Fibre Channel target port you are mapping.

Step 6 Choose the **List** radio button and set the iSCSI initiator node names or IP addresses that you want this virtual iSCSI target to access, or choose the **All** radio button to let the iSCSI target access all iSCSI initiators See the "iSCSI Access Control" section on page 45-27.

Step 7 Choose the **Selected from List** radio button and check each interface you want to advertise the iSCSI targets on or choose the **All** radio button to advertise all interfaces.

Step 8 Click **Apply** to save this change or click **Cancel** to close the dialog box without saving any changes.

LUN Trespass for Storage Port Failover

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

In physical Fibre Channel targets, which are configured to have LUs visible over two Fibre Channel N ports, when the active port fails, the passive port takes over. Some physical Fibre Channel targets require that the trespass feature be used to move the LUs from the active port to the passive port. A statically imported iSCSI target's secondary pWWN option and an additional option of enabling the trespass feature is available for a physical Fibre Channel target with redundant ports. When the active port fails, the passive port becomes active, and if the trespass feature is enabled, the Cisco MDS switch sends a request to the target to move the LUs on the new active port. The iSCSI session switches to use the new active port and the moved LUs are accessed over the new active port (see Figure 45-58).

Figure 45-58 Virtual Target with an Active Primary Port

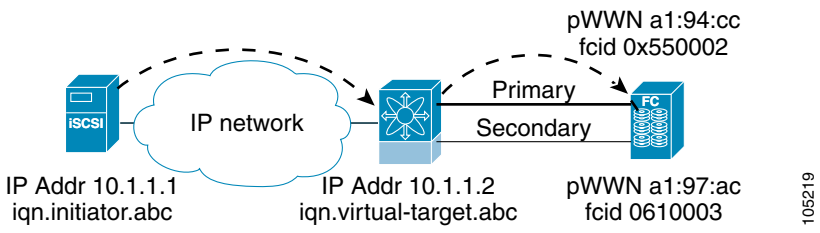

In Device Manager, choose **IP > iSCSI**, select the **Targets** tab, and check the **Trespass Mode** check box to enable the trespass feature for a static iSCSI virtual target.

Multiple IPS Ports Connected to the Same IP Network

Figure 45-59 provides an example of a configuration with multiple Gigabit Ethernet interfaces in the same IP network.

Figure 45-59 Multiple Gigabit Ethernet Interfaces in the Same IP Network

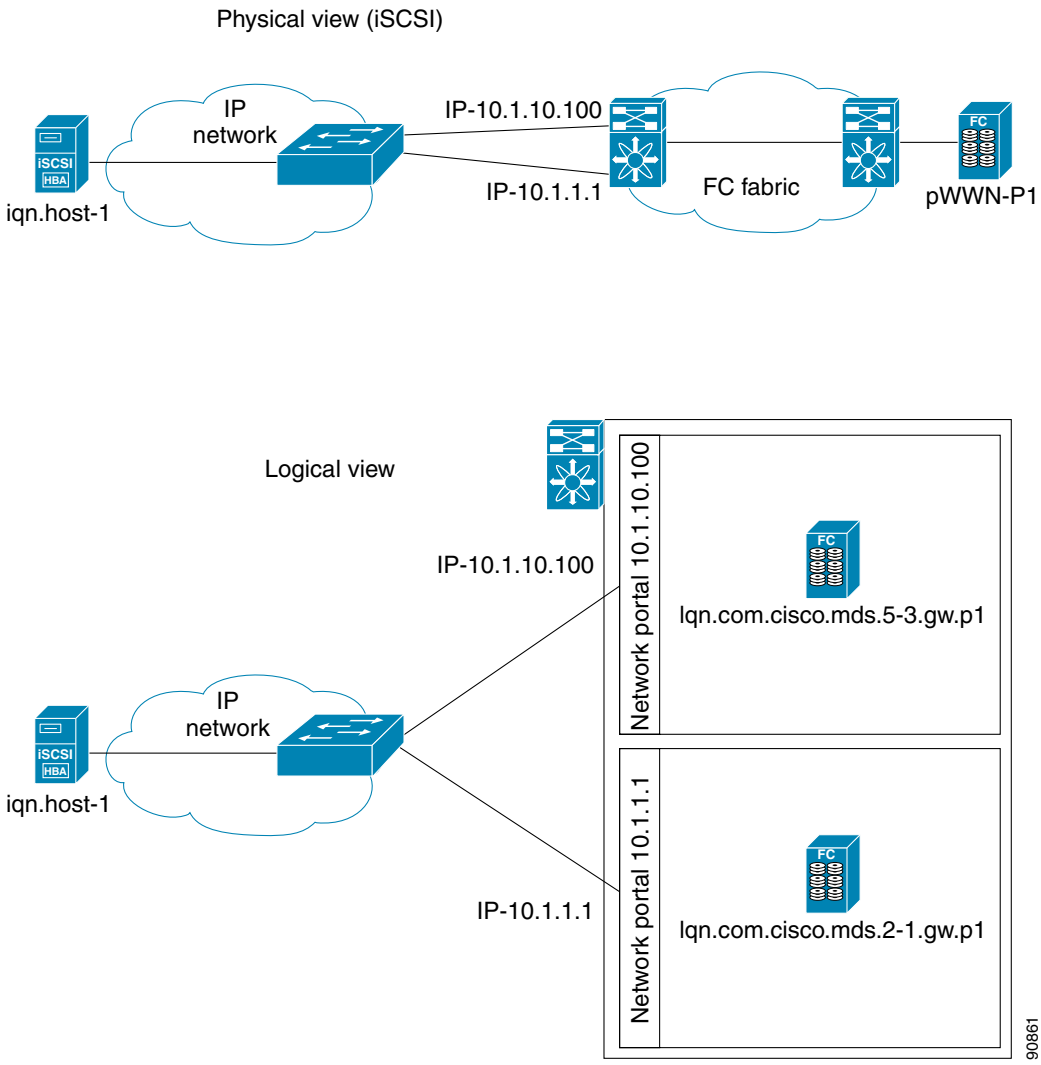

In Figure 45-59, each iSCSI host discovers two iSCSI targets for every physical Fibre Channel target (with different names). The multi-pathing software on the host provides load-balancing over both paths. If one Gigabit Ethernet interface fails, the host multi-pathing software is not affected because it can use the second path.

VRRP-Based High Availability

Figure 45-60 provides an example of a VRRP-based high availability iSCSI configuration.

Figure 45-60 VRRP-Based iSCSI High Availability

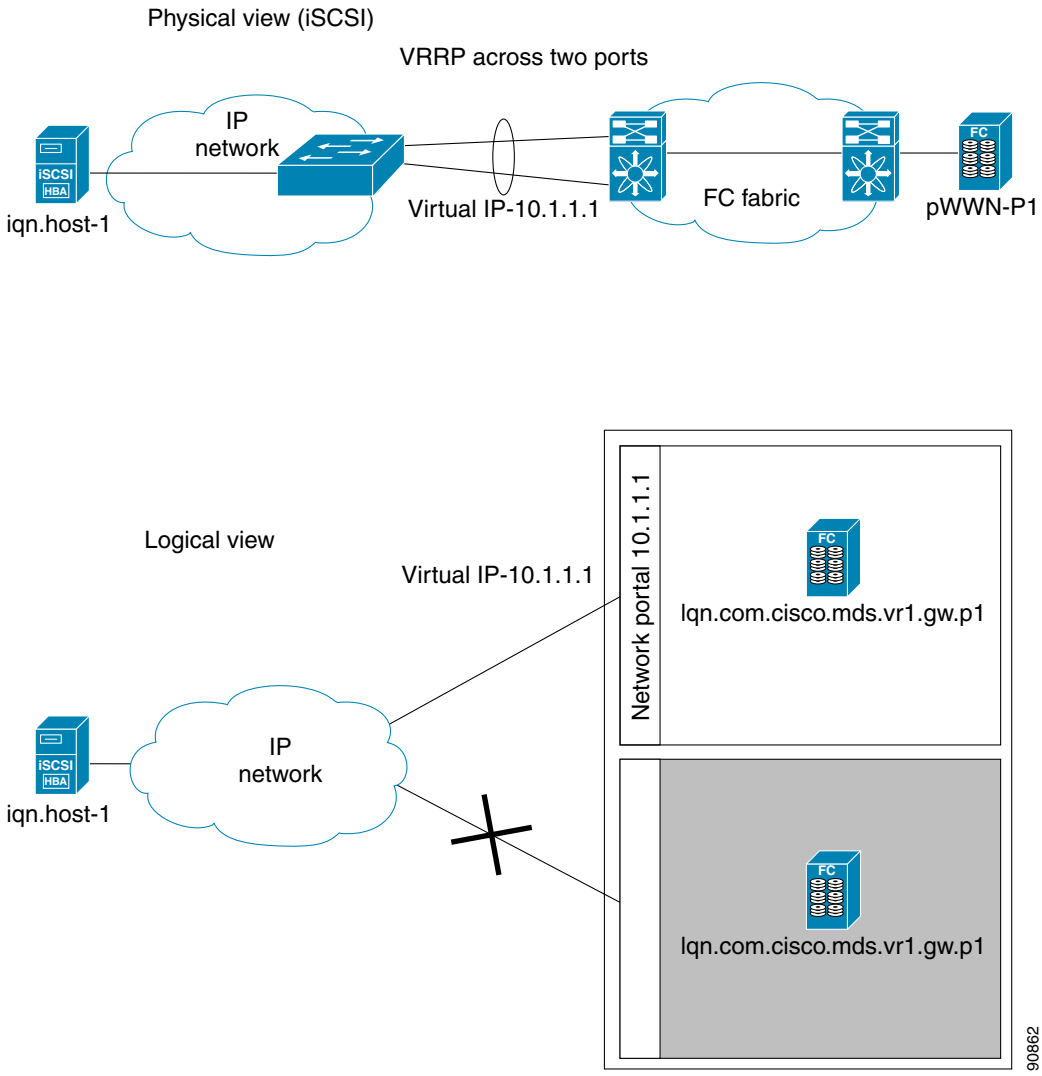

In Figure 45-60, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. When the Gigabit Ethernet interface of the VRRP master fails, the iSCSI session is terminated. The host then reconnects to the target and the session comes up because the second Gigabit Ethernet interface has taken over the virtual IP address as the new master.

Ethernet PortChannel-Based High Availability

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

Figure 45-61 provides a sample Ethernet PortChannel-based high availability iSCSI configuration.

Figure 45-61 Ethernet PortChannel-Based iSCSI High Availability

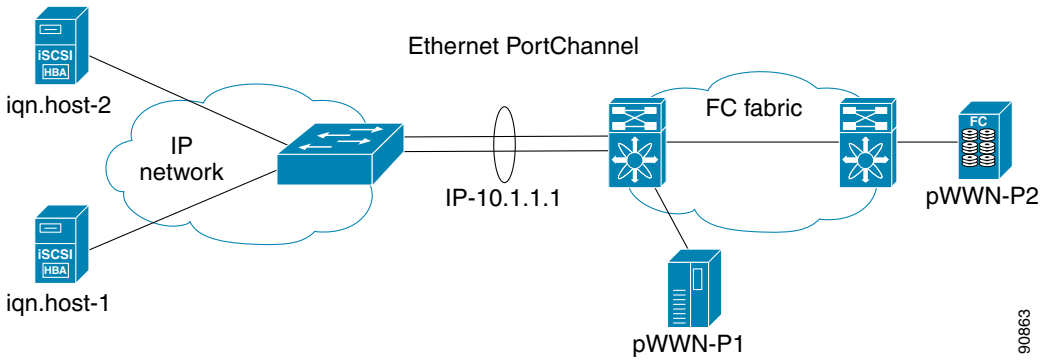

In Figure 45-61, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. The iSCSI session from the iSCSI host to the iSCSI virtual target (on the IPS port) uses one of the two physical interfaces (because an iSCSI session uses one TCP connection). When the Gigabit Ethernet interface fails, the IPS module and the Ethernet switch transparently forwards all the frames on to the second Gigabit Ethernet interface.

iSCSI Authentication Setup Guidelines and Scenarios

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

- No Authentication, page 45-64
- CHAP with Local Password Database, page 45-65
- CHAP with External RADIUS Server, page 45-65
- iSCSI Transparent Mode Initiator, page 45-66
- Target Storage Device Requiring LUN Mapping, page 45-71

Caution Changing the authentication of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the "Changing iSCSI Interface Parameters and the Impact on Load Balancing" section on page 45-52.

No Authentication

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

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

CHAP with Local Password Database

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

Step 1 Set the AAA authentication to use the local password database for the iSCSI protocol:

a. In Fabric Manager, choose **Switches > Security > AAA** in the Physical Attributes pane.

b. Select **Applications** tab in the Information pane.

c. Check the **Local** check box for the iSCSI row and click **Apply Changes**.

Step 2 Set the iSCSI authentication method to require CHAP for all iSCSI clients.:

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

b. Select **Globals** tab in the Information pane.

c. Set the AuthMethod drop-down menu to **chap** and click **Apply Changes**.

Step 3 Configure the user names and passwords for iSCSI users:

a. In Device Manager, choose **Security > iSCSI**.

b. Set the Username, Password and Confirm Password fields.

c. Click **Create** to save these changes or click **Close** to discard any unsaved changes.

Step 4 Verify the global iSCSI authentication setup:

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

b. Select **Globals** tab in the Information pane.

CHAP with External RADIUS Server

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

Step 1 Configure the password for the Cisco MDS switch as RADIUS client to the RADIUS server:

a. In Fabric Manager, choose **Switches > Security > AAA > RADIUS** in the Physical Attributes pane.

b. Select **Default** tab in the Information pane.

c. Set the AuthKey field to the default password and click **Apply Changes**.

Step 2 Configure the RADIUS server IP address:

a. In Fabric Manager, choose **Switches > Security > AAA > RADIUS** in the Physical Attributes pane.

b. Select **Server** tab in the Information pane and click **Create Row**.S

c. Set the Index field to a unique number.

d. Set the IP Type radio button to **ipv4** or **ipv6**.

e. Set the Name or IP Address field to the IP address of the RADIUS server and click **Create**.

Step 3 Create a RADIUS server group and add the RADIUS server to the group:

a. In Fabric Manager, choose **Switches > Security > AAA** in the Physical Attributes pane.

b. Select the **Server Groups** tab in the Information pane and click **Create Row**.

c. Set the Index field to a unique number.I

d. Set the Protocol radio button to **radius**.

e. Set the Name field to the server group name.

f. Set the ServerIDList to the index value of the RADIUS server (as created in Step 2 c.) and click **Create**.

Step 4 Set up the authentication verification for the iSCSI protocol to go to the RADIUS server.

a. In Fabric Manager, choose **Switches > Security > AAA** in the Physical Attributes pane.

b. Select the **Applications** tab in the Information pane.

c. Right-click on the iSCSI row in the Type, SubType, Function column.

d. Set the ServerGroup IDList to the index value of the Server Group (as created in Step 3 c.) and click **Create**.

Step 5 Set up the iSCSI authentication method to require CHAP for all iSCSI clients.

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

b. Select **chap** from the AuthMethod drop-down menu.

c. Click **Apply Changes.**

Step 6 In Fabric Manager, choose **End Devices > iSCSI** in the Physical Attributes pane and select the **Globals** tab in the Information pane to verify that the global iSCSI authentication setup is for CHAP.

Step 7 In Fabric Manager, choose **Switches > Security > AAA** in the Physical Attributes pane and Select the **Applications** tab in the Information pane to verify the AAA authentication information for iSCSI.

To configure an iSCSI RADIUS server, follow these steps:

Step 1 Configure the RADIUS server to allow access from the Cisco MDS switch's management Ethernet IP address.

Step 2 Configure the shared secret for the RADIUS server to authenticate the Cisco MDS switch.

Step 3 Configure the iSCSI users and passwords on the RADIUS server.

iSCSI Transparent Mode Initiator

This scenario assumes the following configuration (see Figure 45-62):

- No LUN mapping or LUN masking or any other access control for hosts on the target device
- No iSCSI login authentication (that is, login authentication set to none)
- The topology is as follows:
  - iSCSI interface 7/1 is configured to identify initiators by IP address.
  - iSCSI interface 7/5 is configured to identify initiators by node name.

- The iSCSI initiator host 1 with IPv4 address 10.11.1.10 and name iqn.1987-05.com.cisco:01.255891611111 connects to IPS port 7/1 is identified using IPv4 address (host 1 = 10.11.1.10).
- The iSCSI initiator host 2 with IPv4 address 10.15.1.10 and node name iqn.1987-05.com.cisco:01.25589167f74c connects to IPS port 7/5.

Figure 45-62 iSCSI Scenario 1

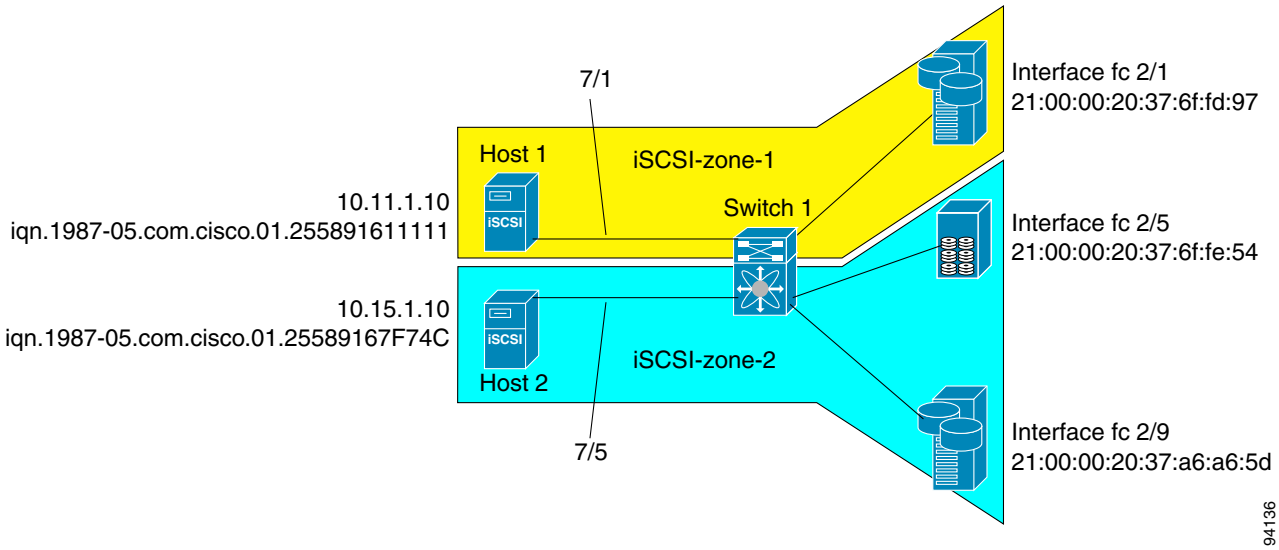

To configure scenario 1 (see Figure 45-62), follow these steps:

Step 1 Configure null authentication for all iSCSI hosts in Cisco MDS switches.

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

b. Select **none** from the AuthMethod drop-down menu in the Information pane.

c. Click **Apply Changes.**

Step 2 Configure iSCSI to dynamically import all Fibre Channel targets into the iSCSI SAN using auto-generated iSCSI target names.

a. In Device Manager, click **IP > iSCSI** and select the **Targets** tab.

b. Check the **Dynamically Import FC Targets** check box.

c. Click **Apply**.

Step 3 Configure the Gigabit Ethernet interface in slot 7 port 1 with an IPv4 address and enable the interface.

a. In Fabric Manager, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.

b. Select the **IP Address** tab in the Information pane and click **Create Row**.

c. Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 1.

d. Click **Create.**

e. Select the General tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 1.

f. Click **Apply Changes.**

Note Host 2 is connected to this port.

Step 4 Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by their IP address, and enable the interface.

a. In Fabric Manager, choose **Switches > Interfaces > FC Logical in** the Physical Attributes pane.

b. Select the **iSCSI** tab in the Information pane.

c. Select **ipaddress** from the Initiator ID Mode drop-down menu and click Apply Changes.

d. In Device Manager, Choose **Interfaces > Ethernet and iSCSI**.

e. Select the **iSCSI** tab.

f. Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 1.

g. Click **Apply**.

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

a. In Fabric Manager, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.

b. Select the **IP Address** tab in the Information pane and click **Create Row**.

c. Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 5.

d. Click **Create.**

e. Select the General tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 5.

f. Click **Apply Changes.**

Step 6 Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by node name and enable the interface.

a. In Fabric Manager, choose **Switches > Interfaces > FC Logical in** the Physical Attributes pane.

b. Select the **iSCSI** tab in the Information pane.

c. Select **name** from the Initiator ID Mode drop-down menu and click Apply Changes.

d. In Device Manager, Choose **Interfaces > Ethernet and iSCSI**.

e. Select the **iSCSI** tab.

f. Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 5.

g. Click **Apply**.

Note Host 1 is connected to this port.

Step 7 Verify the available Fibre Channel targets.

a. In Device Manager, Choose **FC > Name Server**.

b. Select the **General** tab.

Step 8 Create a zone named *iscsi-zone-1* with host 1 and one Fibre Channel target in it.

Note Use the IP address of the host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on IP address.

a. In Fabric Manager, choose **Zones > Edit Local Full Zone Database** from the main menu.
b. Select VSAN 1 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
c. Select the **Zones** folder in the left navigation pane and click **Insert**.
d. Set the Zone Name field to **iscsi-zone-1** and click **OK.**
e. Select the iscsi-zone-1 folder in the left navigation pane and click **Insert**.
f. Set the ZoneBy radio button to**WWN.**
g. Set the Port WWN to the pWWN for the Fibre Channel target (that is, 21:00:00:20:37:6f:fd:97). and click **Add**.
h. Set the ZoneBy radio button to **iSCSI IP Address/Subnet.**
i. Set the IP Address/Mask field to the IP Address for Host 1 iSCSI initiator (10.11.1.10) and click **Add**.

Step 9 Create a zone named *iscsi-zone-2* with host 2 and two Fibre Channel targets in it.

Note Use the symbolic node name of the iSCSI host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on node name.

a. In Fabric Manager, choose **Zones > Edit Local Full Zone Database** from the main menu.
b. Select VSAN 2 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
c. Select the **Zones** folder in the left navigation pane and click **Insert**.
d. Set the Zone Name field to **iscsi-zone-2** and click **OK.**
e. Select the **iscsi-zone-2** folder in the left navigation pane and click **Insert**.
f. Set the ZoneBy radio button to**WWN.**
g. Set the Port WWN to the pWWN for one of the Fibre Channel targets (for example, 21:00:00:20:37:6f:fe:5). and click **Add**.
h. Set the Port WWN to the pWWN for another of the Fibre Channel targets (for example, 21:00:00:20:37:a6:a6:5d). and click **Add**.
i. Set the ZoneBy radio button to **iSCSI name.**
j. Set the Port Name field to the symbolic name for host 2 (iqn.1987-05.com.cisco:01.25589167f74c) and click **Add**.

Step 10 Create a zone set, add the two zones as members, and activate the zone set.

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

a. In Fabric Manager, choose **Zones > Edit Local Full Zone Database** from the main menu.
b. Select VSAN 1 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
c. Select the **Zoneset** folder in the left navigation pane and click **Insert**.
d. Set the Zoneset Name to **zonset-iscsi** and click **OK**.

e. Click on the **zoneset-iscsi** folder and click **Insert**.

f. Set the Zone Name field to **iscsi-zone-1** and click **OK.**

g. Set the Zone Name field to **iscsi-zone-2** and click **OK.**

h. Click **Activate** to activate the new zone set.

i. Click **Continue Activation** to finish the activation.

Step 11 Bring up the iSCSI hosts (host 1 and host 2).

Step 12 Show all the iSCSI sessions.

a. In Device Manager, choose **Interfaces > Monitor > Ethernet** and select the **iSCSI Sessions** tab to show all the iSCSI sessions.

b. In Device Manager, choose **IP > iSCSI** and select the **Session Initiators** tab.

c. Click **Details.**

Step 13 In Fabric Manager, choose **End Devices > iSCSI** in the Physical Attributes pane to verify the details of the two iSCSI initiators

Step 14 In Fabric Manager, choose **Zones > Edit Local Full Zone Database** from the main menu to view the active zone set. The iSCSI initiators' FC IDs are resolved.

Step 15 In Device Manager, Choose **FC > Name Server.** The Fibre Channel name server shows the virtual N ports created for the iSCSI hosts.

Step 16 In Device Manager, Choose **FC > Name Server** and select the **Advanced** tab**.** Verify the detailed output of the iSCSI initiator nodes in the Fibre Channel name server.

Target Storage Device Requiring LUN Mapping

Sample scenario 2 assumes the following configuration (see Figure 45-63):

- Access control is based on Fibre Channel zoning.
- There is target-based LUN mapping or LUN masking.
- There is no iSCSI authentication (none).
- The iSCSI initiator is assigned to different VSANs.

Figure 45-63 iSCSI Scenario 2

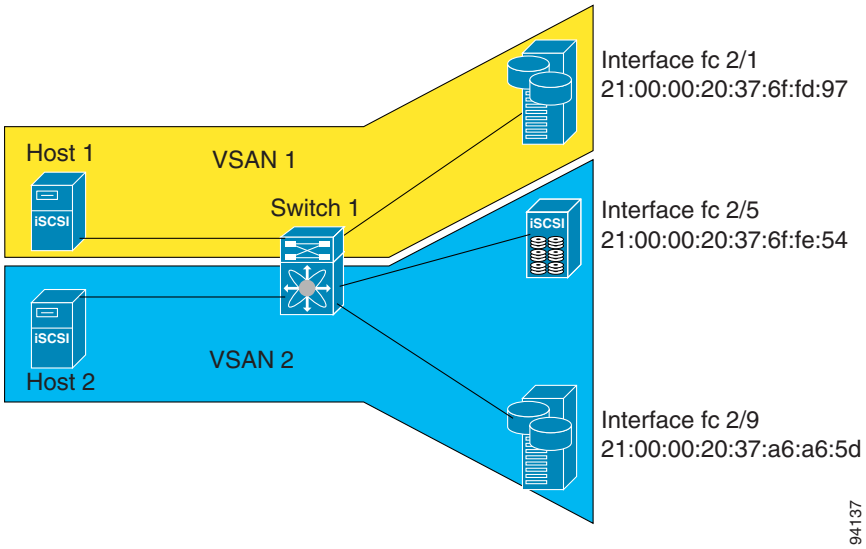

To configure scenario 2 (see Figure 45-63), follow these steps:

Step 1 Configure null authentication for all iSCSI hosts.

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

b. Select **none** from the AuthMethod drop-down menu in the Information pane.

c. Click **Apply Changes.**

Step 2 Configure iSCSI to dynamically import all Fibre Channel targets into the iSCSI SAN using auto-generated iSCSI target names.

a. In Device Manager, click **IP > iSCSI** and select the **Targets** tab.

b. Check the **Dynamically Import FC Targets** check box.

c. Click **Apply**.

Step 3 Configure the Gigabit Ethernet interface in slot 7 port 1 with an IPv4 address and enable the interface.

a. In Fabric Manager, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.

b. Select the **IP Address** tab in the Information pane and click **Create Row**.

c. Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 1.

d. Click **Create.**

e. Select the **General** tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 1.

f. Click **Apply Changes.**

Step 4 Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by their IP address and enable the interface.

a. In Fabric Manager, choose **Switches > Interfaces > FC Logical in** the Physical Attributes pane.

b. Select the **iSCSI** tab in the Information pane.

c. Select **ipaddress** from the Initiator ID Mode drop-down menu and click Apply Changes.

d. In Device Manager, Choose **Interfaces > Ethernet and iSCSI**.

e. Select the **iSCSI** tab.

f. Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 1.

g. Click **Apply**.

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

a. In Fabric Manager, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.

b. Select the **IP Address** tab in the Information pane and click **Create Row**.

c. Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 5.

d. Click **Create.**

e. Select the **General** tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 5.

f. Click **Apply Changes.**

Step 6 Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by IP address and enable the interface.

a. In Fabric Manager, choose **Switches > Interfaces > FC Logical in** the Physical Attributes pane.

b. Select the **iSCSI** tab in the Information pane.

c. Select **ipaddress** from the Initiator ID Mode drop-down menu and click Apply Changes.

d. In Device Manager, Choose **Interfaces > Ethernet and iSCSI**.

e. Select the **iSCSI** tab.

f. Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 5.

g. Click **Apply**.

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

h. In Device Manager, Choose **IP > iSCSI**.

i. Select the **Initiators** tab.

j. Check the **Node Address Persisten**t and **Node Address System-assigned** check boxes the Host 1 iSCSI initiator.

k. Click **Apply**.

Step 8 Configure for static pWWN for Host 2.

a. In Device Manager, Choose **IP > iSCSI**.

b. Select the **Initiators** tab.

c. Right-click on the Host 2 iSCSI initiator and click Edit pWWN.

d. Select **1** from the System-assigned Num field and click **Apply**.

Step 9 View the configured WWNs.

Note The WWNs are assigned by the system. The initiators are members of different VSANs.

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

b. Select the **Initiators** tab.

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

Note Use the IP address of the host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on IP address.

a. In Fabric Manager, choose **Zones > Edit Local Full Zone Database** from the main menu.

b. Select VSAN 1 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.

c. Select the **Zones** folder in the left navigation pane and click **Insert**.

d. Set the Zone Name field to **iscsi-zone-1** and click **OK.**

e. Select the iscsi-zone-1 folder in the left navigation pane and click **Insert**.

f. Set the ZoneBy radio button to**WWN.**

g. Set the Port WWN to the pWWN for the Fibre Channel target (that is, 21:00:00:20:37:6f:fd:97). and click **Add**.

h. Set the ZoneBy radio button to **iSCSI IP Address/Subnet.**

i. Set the IP Address/Mask field to the IP Address for Host 1 iSCSI initiator (10.11.1.10) and click **Add**.

Note Fibre Channel storage for zone membership for the iSCSI initiator, either the iSCSI symbolic node name or the pWWN, can be used. In this case, the pWWN is persistent.

Step 11 Create a zone set in VSAN 1 and activate it.

a. In Fabric Manager, choose **Zones > Edit Local Full Zone Database** from the main menu.

b. Select VSAN 1 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.

c. Select the **Zoneset** folder in the left navigation pane and click **Insert**.

d. Set the Zoneset Name to **zonset-iscsi-1** and click **OK**.

e. Click on the **zoneset-iscsi-1** folder and click **Insert**.

f. Set the Zone Name field to **iscsi-zone-1** and click **OK.**

g. Click **Activate** to activate the new zone set.

h. Click **Continue Activation** to finish the activation.

Step 12 Create a zone with host 2 and two Fibre Channel targets.

Note If the host is in VSAN 2, the Fibre Channel targets and zone must also be in VSAN 2.

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

a. In Fabric Manager, choose **Zones > Edit Local Full Zone Database** from the main menu.
b. Select VSAN 2 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
c. Select the **Zones** folder in the left navigation pane and click **Insert**.
d. Set the Zone Name field to **iscsi-zone-2** and click **OK.**
e. Select the **iscsi-zone-2** folder in the left navigation pane and click **Insert**.
f. Set the ZoneBy radio button to**WWN.**
g. Set the Port WWN to the pWWN for one of the Fibre Channel targets (for example, 21:00:00:20:37:6f:fe:5). and click **Add**.
h. Set the Port WWN to the pWWN for another of the Fibre Channel targets (for example, 21:00:00:20:37:a6:a6:5d). and click **Add**.
i. Set the ZoneBy radio button to **iSCSI IP Address/Subnet.**
j. Set the IP Address/Mask field to the IP Address for Host 2 iSCSI initiator (10.15.1.11) and click **Add**.

Step 13 Create a zone set in VSAN 2 and activate it.

a. In Fabric Manager, choose **Zones > Edit Local Full Zone Database** from the main menu.
b. Select VSAN 2 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
c. Select the **Zoneset** folder in the left navigation pane and click **Insert**.
d. Set the Zoneset Name to **zonset-iscsi-2** and click **OK**.
e. Click on the **zoneset-iscsi-2** folder and click **Insert**.
f. Set the Zone Name field to **iscsi-zone-2** and click **OK.**
g. Click **Activate** to activate the new zone set.
h. Click **Continue Activation** to finish the activation.

Step 14 Start the iSCSI clients on both hosts

Step 15 Show all the iSCSI sessions.

a. In Device Manager, choose **Interfaces > Monitor > Ethernet** and select the **iSCSI Sessions** tab to show all the iSCSI sessions.
b. In Device Manager, choose **IP > iSCSI** and select the **Session Initiators** tab.
c. Click **Details.**

Step 16 In Fabric Manager, choose **End Devices > iSCSI** in the Physical Attributes pane to verify the details of the two iSCSI initiators.

Step 17 In Fabric Manager, choose **Zones > Edit Local Full Zone Database** from the main menu to view the active zone set. The iSCSI initiators' FC IDs are resolved

Step 18 In Device Manager, Choose **FC > Name Server.** The Fibre Channel name server shows the virtual N ports created for the iSCSI hosts.

Step 19 In Device Manager, Choose **FC > Name Server** and select the **Advanced** tab**.** Verify the detailed output of the iSCSI initiator nodes in the Fibre Channel name server.

iSNS

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

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

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

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

This section includes the following topics:

About iSNS Client Functionality

The iSNS client functionality on each IPS interface (Gigabit Ethernet interface or subinterface or PortChannel) registers information with an iSNS server. You must specify an iSNS server's IP address by creating an iSNS profile, adding the server's IP address to it, and then assigning (or "tagging") the profile to the interface. An iSNS profile can be tagged to one or more interfaces.

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

Statically mapped virtual targets are registered if the associated Fibre Channel pWWN is present in the FCNS database and no access control configuration prevents it. A dynamically mapped target is registered if dynamic target importing is enabled. See the "Presenting Fibre Channel Targets as iSCSI Targets" section on page 45-7 for more details on how iSCSI imports Fibre Channel targets.

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

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

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

Creating an iSNS Client Profile

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

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

You see the iSCSI Configuration in the Information pane.

Step 2 Select the **iSNS** tab. You see the iSNS profiles configured (see Figure 45-64).

Figure 45-64 iSNS Profiles in Fabric Manager

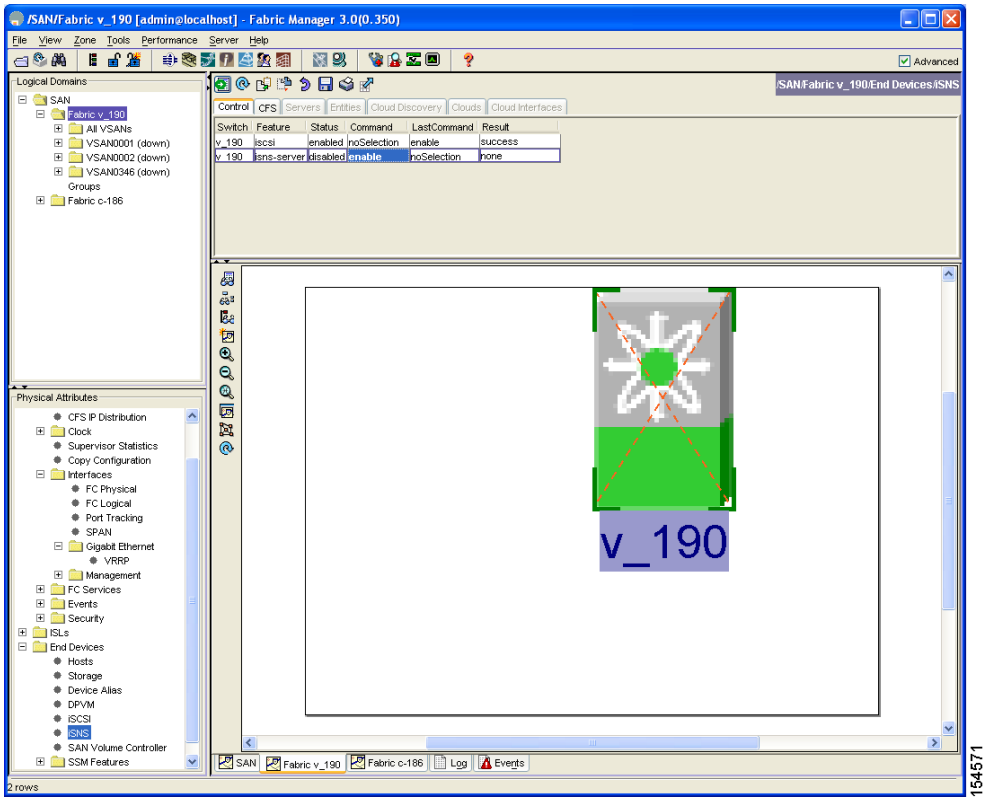

Step 3 Click **Create Row**.

You see the Create iSNS Profiles dialog box.

Step 4 Set the ProfileName field to the iSNS profile name that you want to create.

Step 5 Set the ProfileAddr field to the IP address of the iSNS server.

Step 6 Click **Create** to save these changes or click **Close** to discard any unsaved changes.

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

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

You see the iSCSI Configuration in the Information pane.

Step 2 Select the **iSNS** tab.

You see the iSNS profiles configured (see Figure 45-64).

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

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

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

You see the Gigabit Ethernet Configuration in the Information pane.

Step 2 Select the **iSNS** tab.

You see the iSNS profiles configured for these interfaces (see Figure 45-65).

Figure 45-65 iSNS Profiles in Fabric Manager

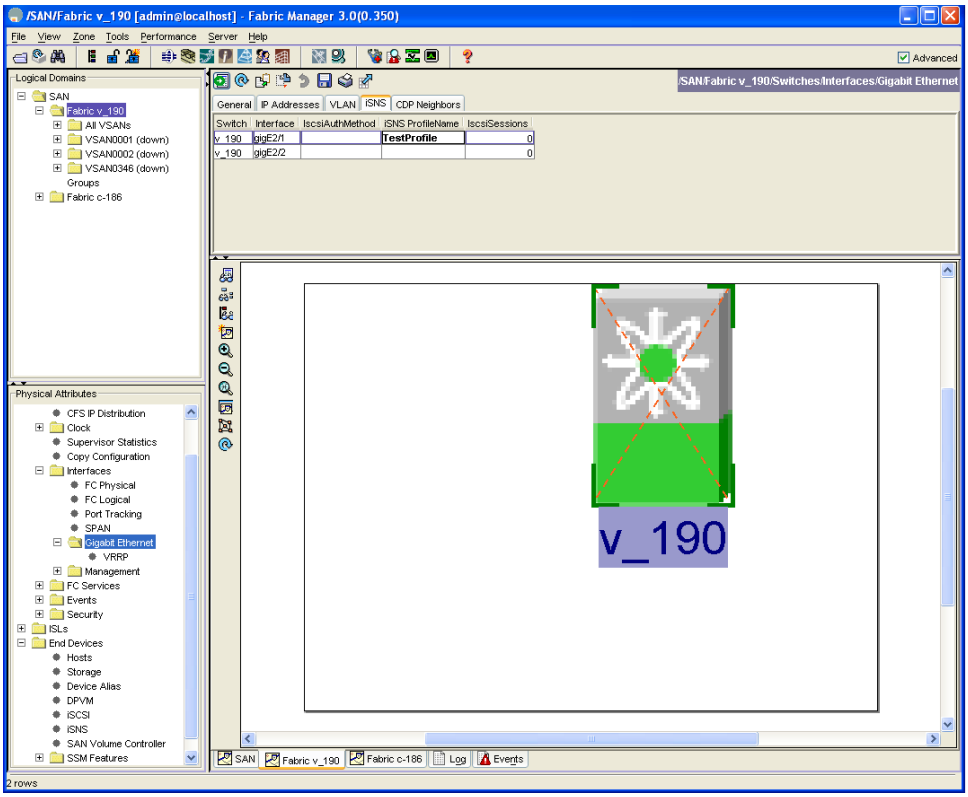

Step 3 Set the iSNS ProfileName field to the iSNS profile name that you want to add to this interface.

Step 4 Click the **Apply Changes** icon to save these changes or click the **Undo Changes** icon to discard any unsaved changes.

To untag a profile to an interface using Fabric Manage, follow these steps:

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

You see the Gigabit Ethernet Configuration in the Information pane.

Step 2 Select the **iSNS** tab.

You see the iSNS profiles configured for these interfaces.

Step 3 Right-click on iSNS ProfileName field that you want to untag and delete the text in that field.

Step 4 Click the **Apply Changes** icon to save these changes or click the **Undo Changes** icon to discard any unsaved changes.

About iSNS Server Functionality

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

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

Example Scenario

The iSNS server provides uniform access control across Fibre Channel and iSCSI devices by utilizing both Fibre Channel zoning information and iSCSI access control information and configuration. An iSCSI initiator acting as an iSNS client only discovers devices it is allowed to access based on both sets of access control information. Figure 45-66 provides an example of this scenario.

Figure 45-66 Using iSNS Servers in the Cisco MDS Environment

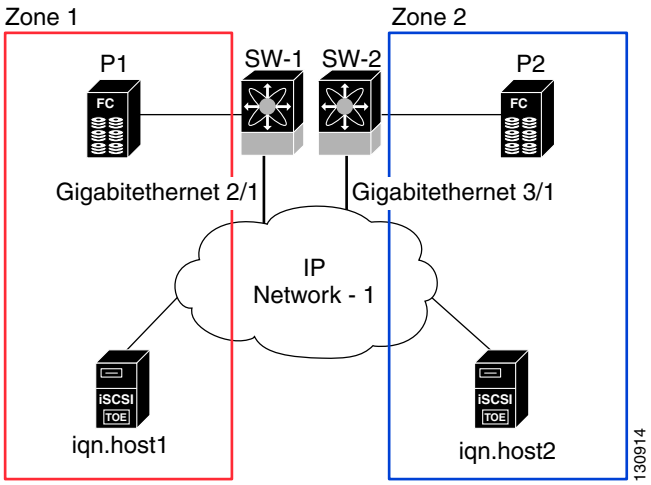

In Figure 45-66, iqn.host1 and iqn.host2 are iSCSI initiators. P1 and P2 are Fibre Channel targets. The two initiators are in different zones: Zone 1 consists of iqn.host1 and target P1, and Zone 2 consists of iqn.host2 and target P2. iSNS server functionality is enabled on both switches, SW-1 and SW-2. The registration process proceeds as follows:

1. Initiator iqn.host1 registers with SW-1, port Gigabitethernet2/1.
2. Initiator iqn.host2 registers with SW-2, port Gigabitethernet3/1.
3. Initiator iqn.host1 issues an iSNS query to SW-1 to determine all accessible targets.
4. The iSNS server in turn queries the Fibre Channel name server (FCNS) to obtain a list of devices that are accessible (that is, in the same zone) by the query originator. This query yields only P1.
5. The iSNS server then queries its own database to convert the Fibre Channel devices to the corresponding iSCSI targets. This is based on the iSCSI configuration, such as virtual-target and its access control setting or whether the dynamic Fibre Channel target import feature is enabled or disabled.
6. The iSNS server sends a response back to the query initiator. This response contains a list all iSCSI portals known to the iSNS server. This means iqn.host1 can choose to log in to target P1 through either SW-1 (at Gigabitethernet 2/1) or SW-2 (at Gigabitethernet 3/1).
7. If the initiator chooses to log in to SW-1 and later that port becomes inaccessible (for example, Gigabitethernet 2/1 goes down), the initiator has the choice to move to connect to target P1 through port Gigabitethernet 3/1 on SW-2 instead.
8. If the target either goes down or is removed from the zone, the iSNS server sends out an iSNS State Change Notification (SCN) message to the initiator so that the initiator can remove the session.

Configuring iSNS Servers

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

This section includes the following topics:

Enabling the iSNS Server

Before the iSNS server feature can be enabled, iSCSI must be enabled (see the "Enabling iSCSI" section on page 45-4). When you disable iSCSI, iSNS is automatically disabled. When the iSNS server is enabled on a switch, every IPS port whose corresponding iSCSI interface is up is capable of servicing iSNS registration and query requests from external iSNS clients.

To enable the iSNS server using Fabric Manager, follow these steps:

Step 1 Choose **End Devices > iSNS**. You see the iSNS configuration in the Information pane.

Step 2 Select the **Control** tab and select **enable** from the **Command** drop-down menu for the iSNS server feature.

Step 3 Click **Apply Changes** to save this change or click **Undo Changes** to discard any unsaved changes.

Note If you are using VRRP IP address for discovering targets from iSNS client, ensure that the IP address is created using the secondary option.

iSNS Configuration Distribution

You can use the CFS infrastructure to distribute the iSCSI initiator configuration to iSNS servers across the fabric. This allows the iSNS server running on any switch to provide a querying iSNS client a list of iSCSI devices available anywhere on the fabric. For information on CFS, see Chapter 12, "Using the CFS Infrastructure."

To enable iSNS configuration distribution using Fabric Manager, follow these steps:

Step 1 Choose **End Devices > iSNS**. You see the iSNS configuration in the Information pane.

Step 2 Select the **CFS** tab and select **enable** from the Admin drop-down menu for iSNS.

Step 3 Select **enable** from the Global drop-down menu for iSNS.

Step 4 Click **Apply Changes** to save this change or click **Undo Changes** to discard any unsaved changes.

Configuring the ESI Retry Count

The iSNS client registers information with its configured iSNS server using an iSNS profile. At registration, the client can indicate an entity status inquiry (ESI) interval of 60 seconds or more. If the client registers with an ESI interval set to zero (0), then the server does not monitor the client using ESI. In such cases, the client's registrations remain valid until explicitly deregistered or the iSNS server feature is disabled.

The ESI retry count is the number of times the iSNS server queries iSNS clients for their entity status. The default ESI retry count is 3. The client sends the server a response to indicate that it is still alive. If the client fails to respond after the configured number of retries, the client is deregistered from the server.

To configure the ESI retry count for an iSNS server using Fabric Manager, follow these steps:

Step 1 Choose **End Devices > iSNS**. You see the iSNS configuration in the Information pane.

Step 2 Select the **Servers** tab. You see the configured iSNS servers.

Step 3 Set the **ESI NonResponse Threshold** field to the ESI retry count value.

Step 4 Click **Apply Changes** to save this change or click **Undo Changes** to discard any unsaved changes.

iSNS Client Registration and Deregistration

An iSNS client cannot query the iSNS server until it has registered.

iSNS client deregistration can occur either explicitly or when the iSNS server detects that it can no longer reach the client (through ESI monitoring).

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

Target Discovery

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

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

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

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

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

iSNS Cloud Discovery

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

About Cloud Discovery

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

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

Cloud discovery is initiated when the following events occur:

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

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

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

Configuring iSNS Cloud Discovery

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

Enabling iSNS Cloud Discovery

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

Step 1 Choose **End Devices > iSNS**. You see the iSNS configuration in the Information pane.

Step 2 Select the **Control** tab and select **enable** from the **Command** drop-down menu for the cloud discovery feature.

Step 3 Click **Apply Changes** to save this change or click **Undo Changes** to discard any unsaved changes.

Initiating On-Demand iSNS Cloud Discovery

To initiate on-demand iSNS cloud discovery using Fabric Manager, follow these steps:

Step 1 Choose **End Devices > iSNS**. You see the iSNS configuration in the Information pane.

Step 2 Select the **Cloud Discovery** tab and check the **Manual Discovery** check box.

Step 3 Click **Apply Changes** to save this change or click **Undo Changes** to discard any unsaved changes.

Configuring Automatic iSNS Cloud Discovery

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

Step 1 Choose **End Devices > iSNS**. You see the iSNS configuration in the Information pane.

Step 2 Select the **Cloud Discovery** tab and check the **AutoDiscovery** check box.

Step 3 Click **Apply Changes** to save this change or click **Undo Changes** to discard any unsaved changes.

Configuring iSNS Cloud Discovery Distribution

To configure iSNS cloud discovery CFS distribution using Fabric Manager, follow these steps:

Step 1 Choose **End Devices > iSNS**. You see the iSNS configuration in the Information pane.

Step 2 Select the **CFS** tab and select **enable** from the Admin drop-down menu for the cloud discovery feature.

Step 3 Select **enable** from the Global drop-down menu for the cloud discovery feature.

Step 4 Click **Apply Changes** to save this change or click **Undo Changes** to discard any unsaved changes.

Default Settings

Table 45-2 lists the default settings for iSCSI parameters.

Table 45-2 Default iSCSI Parameters

Parameters	Default
Number of TCP connections	One per iSCSI session.
minimum-retransmit-time	300 msec.
keepalive-timeout	60 seconds.
max-retransmissions	4 retransmissions.
PMTU discovery	Enabled.
pmtu-enable reset-timeout	3600 sec.
SACK	Enabled.
max-bandwidth	1Gbps
min-available-bandwidth	70 Mbps.
round-trip-time	1 msec.
Buffer size	4096 KB.
Control TCP and data connection	No packets are transmitted.
TCP congestion window monitoring	Enabled.
Burst size	50 KB.
Jitter	500 microseconds.
TCP connection mode	Active mode is enabled.
Fibre Channel targets to iSCSI	Not imported.
Advertising iSCSI target	Advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces.
iSCSI hosts mapping to virtual Fibre Channel hosts	Dynamic mapping.
Dynamic iSCSI initiators	Members of the VSAN 1.
Identifying initiators	iSCSI node names.
Advertising static virtual targets	No initiators are allowed to access a virtual target (unless explicitly configured).
iSCSI login authentication	CHAP or none authentication mechanism.
revert-primary-port	Disabled.
Header and data digest	Enabled automatically when iSCSI initiators send requests. This feature cannot be configured and is not available in store-and-forward mode.
iSNS registration interval	60 sec (not configurable).
iSNS registration interval retries	3.
Fabric distribution	Disabled.

Table 45-3 lists the default settings for iSLB parameters.

Table 45-3 Default iSLB Parameters

Parameters	Default
Fabric distribution	Disabled.
Load balancing metric	1000.

CHAPTER 46

Configuring IP Services

Cisco MDS 9000 Family switches can route IP traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature is used to route traffic between VSANs. To do so, each VSAN must be in a different IP subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMS):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding on the in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

Switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. VRRP is a restartable application that provides a redundant, alternate path to the gateway switch.

Note For information about configuring IPv6, see Chapter 48, “Configuring IPv6 for Gigabit Ethernet Interfaces.”

This chapter includes the following sections:

Traffic Management Services

In-band options are compliant with and use the RFC 2625 standards. An NMS host running the IP protocol over an FC interface can access the switch using the IPFC functionality. If the NMS does not have a Fibre Channel HBA, in-band management can still be performed using one of the switches as an access point to the fabric (see Figure 46-1).

Figure 46-1 Management Access to Switches

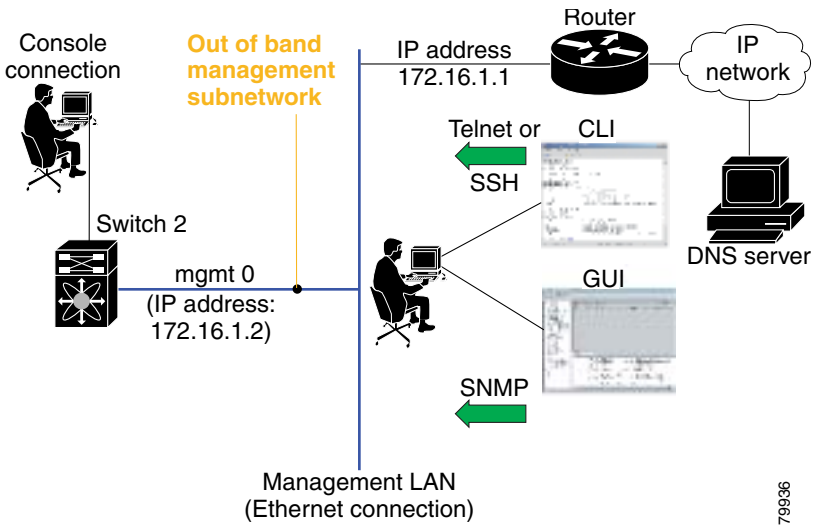

Management Interface Configuration

The management interface on the switch allows multiple simultaneous Telnet or SNMP sessions. You can remotely configure the switch through the management interface, but first you must configure IP version 4 parameters (IP address, subnet mask) or an IP version 6 (IPv6) address and prefix length so that the switch is reachable. For information on configuring IPv6 addresses, see Chapter 48, "Configuring IPv6 for Gigabit Ethernet Interfaces."

On director class switches, a single IP address is used to manage the switch. The active supervisor module's management (mgmt0) interface uses this IP address. The mgmt0 interface on the standby supervisor module remains in an inactive state and cannot be accessed until a switchover happens. After a switchover, the mgmt0 interface on the standby supervisor module becomes active and assumes the same IP address as the previously active supervisor module.

Note The port on the Ethernet switch to which the MDS management interface is connected should be configured as a host port (also known as access port) instead of a switch port. Spanning tree configuration for that port (on the Ethernet switch) should disabled. This helps avoid the delay in the

management port coming up due to delay from Ethernet spanning tree processing that the Ethernet switch would run if enabled. FoR Cisco Ethernet switches, use either the **switchport host** command in IOS is or the **set port host** in Catalyst OS. Refer to the configuration guide for your Ethernet switch.

Note Before you begin to configure the management interface manually, obtain the switch's IP address and IP subnet mask. Also make sure the console cable is connected to the console port.

To configure the mgmt0 Ethernet interface using Device Manager, follow these steps:

Step 1 Select **Interface > Mgmt > Mgmt0**.

Step 2 Enter the description.

Step 3 Select the administrative state of the interface.

Step 4 Check the **CDP** check box to enable CDP.

Step 5 Enter the IP address mask.

Step 6 Click **Apply** to apply the changes.

Default Gateway

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

This section includes the following topics:

About the Default Gateway

The default gateway IPv4 address should be configured along with the IPv4 static routing attributes (IP default network, destination prefix, and destination mask, and next hop address).

Tip If you configure the static route IP forwarding and the default-network details, these IPv4 addresses will be used regardless of the default-gateway being enabled or disabled. If these IP addresses are configured but not available, the switch will fall back to using the default gateway IP address, if you have configured it. Be sure to configure IP addresses for all entries in the switch.

See the "Initial Setup Routine" section on page 2-2 for more information on configuring the IP addresses for all entries in the switch.

Configuring the Default Gateway

To configure an IP route or identify the default gateway using Device Manager, follow these steps:

Step 1 Choose **IP > Routes**.

You see the IP Routes window.

Step 2 Create a new IP route or identify the default gateway on a switch by clicking **Create**.

You see the Create IP Routes window.

Step 3 Complete the fields in this window. Configure a static route, by entering the destination network ID and subnet mask in the Dest and Mask fields. Configure a default gateway by entering the IP address of the seed switch in the Gateway field.

Step 4 Click **Create** to add the IP route.

IPv4 Default Network Configuration

If you assign the IPv4 default network address, the switch considers routes to that network as the last resort. If the IPv4 default network address is not available, the switch uses the IPv4 default gateway address. For every network configured with the IPv4 default network address, the switch flags that route as a candidate default route, if the route is available.

Tip If you configure the static route IP forwarding and the default network details, these IPv4 addresses will be used regardless of the default gateway being enabled or disabled. If these IPv4 addresses are configured and not available, the switch will fall back to using the default gateway IPv4 address, if you have configured it. Be sure to configure IPv4 addresses for all entries in the switch if you are using IPv4.

See the "Initial Setup Routine" section on page 2-2 for more information on configuring the IP addresses for all entries in the switch.

When the Ethernet interface is configured, the switch should point to the gateway router for the IP network. The host accesses the gateway using a gateway switch. This gateway switch is configured as the default gateway. The other switches in the fabric that are connected to the same VSAN as the gateway switch can also be connected through the gateway switch. Every interface connected to this VSAN should be configured with the VSAN IPv4 address of the gateway switch (see "IPv4 Default Network Configuration" section on page 46-4).

Figure 46-2 Overlay VSAN Functionality

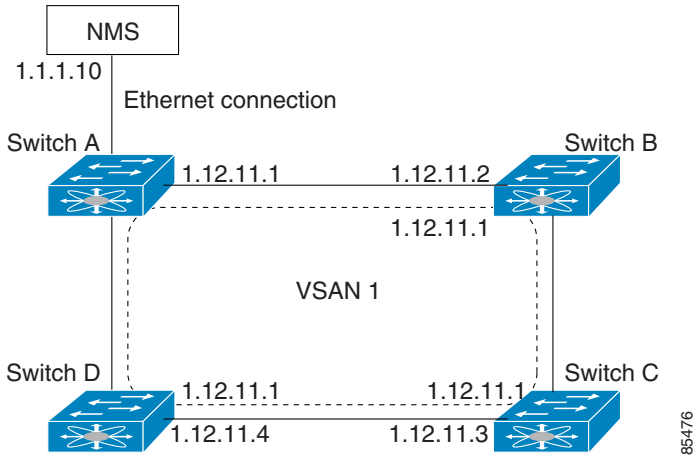

In Figure 46-2, switch A has the IPv4 address 1.12.11.1, switch B has the IPv4 address 1.12.11.2, switch C has the IPv4 address 1.12.11.3, and switch D has the IPv4 address 1.12.11.4. Switch A is the gateway switch with the Ethernet connection. The NMS uses the IPv4 address 1.1.1.10 to connect to the gateway switch. Frames forwarded to any switch in the overlaid VSAN 1 are routed through the gateway switch. Configuring the gateway switch's IPv4 address (1.12.11.1) in the other switches enable the gateway switch to forward the frame to the intended destination. Similarly, if a non-gateway switch in the VSAN forwards a frame to the Ethernet world, the frame is routed through the gateway switch.

When forwarding is disabled (default), IP frames are not sent from one interface to another. In these cases, the software performs local IP routing between two switches using the in-band option for Fibre Channel traffic and the mgmt0 option for Ethernet traffic.

When a VSAN is created, a VSAN interface is not created automatically. You need to specifically create the interface (see the "VSAN Interfaces" section on page 18-16).

IPFC

Once the VSAN interface is created, you can specify the IP address for that VSAN. You can assign an IPv4 address or an IPv6 address.

Note See the Chapter 48, "Configuring IPv6 for Gigabit Ethernet Interfaces" for information about configuring IPv6 on the Cisco MDS 9000 Family switches.

IPv4 Static Routes

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

Note For information about IPv6 static routing, see the "Configuring IPv6 Static Routes" section on page 48-23.

Static routing is a mechanism to configure IPv4 routes on the switch. You can configure more than one static route.

If a VSAN has multiple exit points, configure static routes to direct traffic to the appropriate gateway switch. IPv4 routing is disabled by default on any gateway switch between the out-of-band management interface and the default VSAN, or between directly connected VSANs.

Overlay VSANs

This section describes overlay VSANs and how to configure them.

This section includes the following topics:

About Overlay VSANs

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

Configuring Overlay VSANs

To configure an overlay VSAN, follow these steps:

Step 1 Add the VSAN to the VSAN database on all switch in the fabric.

Step 2 Create a VSAN interface for the VSAN on all switches in the fabric. Any VSAN interface belonging to the VSAN has an IP address in the same subnet. Create a route to the IPFC cloud on the IP side.

Step 3 Configure a default route on every switch in the Fibre Channel fabric pointing to the switch that provides NMS access.

Step 4 Configure the default gateway (route) and the IPv4 address on switches that point to the NMS (see Figure 46-3).

Figure 46-3 Overlay VSAN Configuration Example

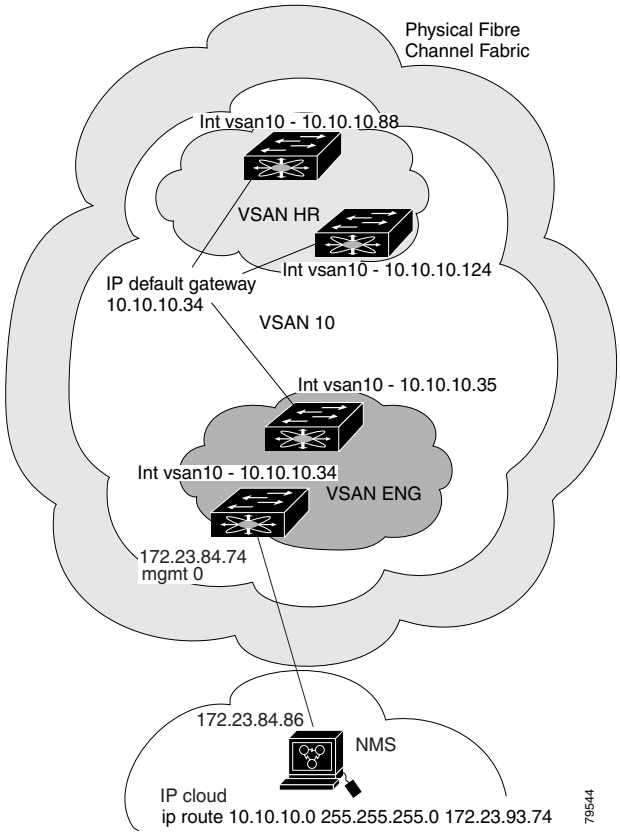

Note To configure the management interface displayed in Figure 46-3, set the default gateway to an IPv4 address on the Ethernet network.

Multiple VSAN Configuration

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

To configure multiple VSANs, follow these steps:

Step 1 Add the VSAN to the VSAN database on any switch in the fabric.

Step 2 Create a VSAN interface for the appropriate VSAN on any switch in the fabric.

Step 3 Assign an IP address on every VSAN interface on the same subnet as the corresponding VSAN.

Step 4 Define the multiple static routes on the Fibre Channel switches and the IP cloud (see Figure 46-4).

Figure 46-4 Multiple VSAN Configuration Example

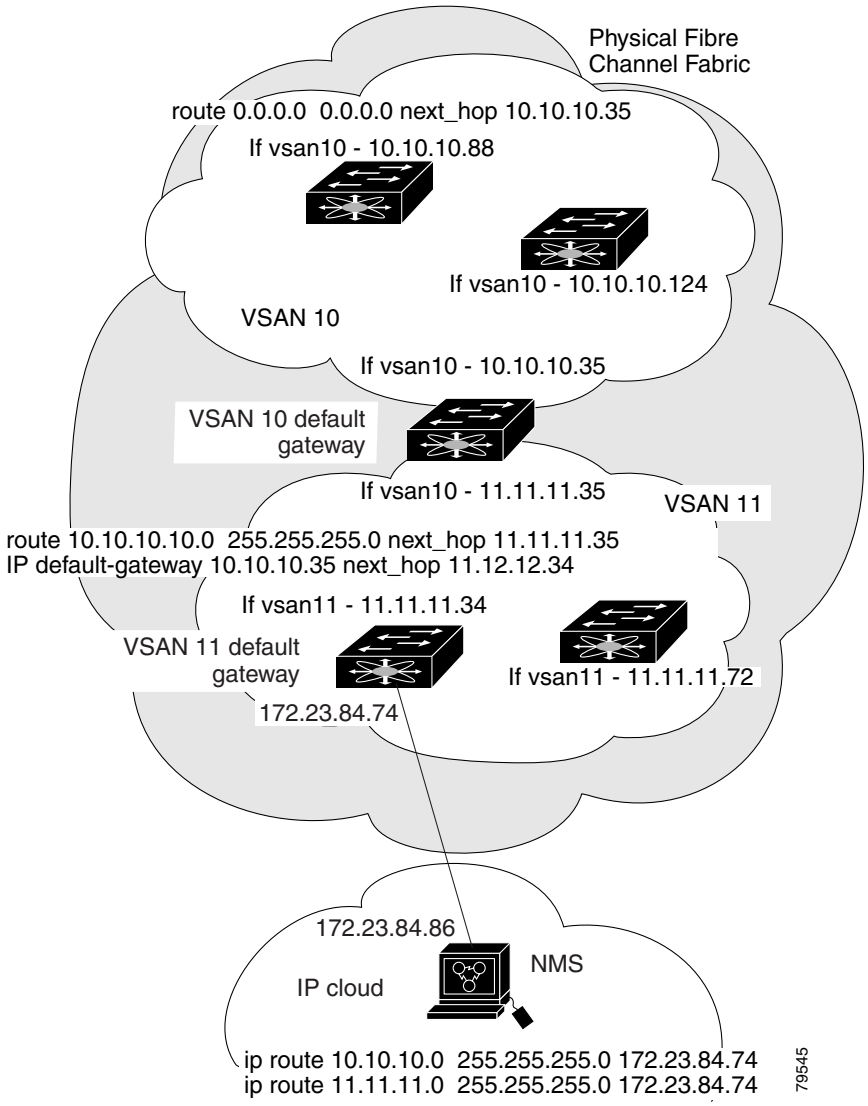

Virtual Router Redundancy Protocol

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

This section includes the following topics:

About VRRP

VRRP provides a redundant alternative path to the gateway switch, which has connectivity to the NMS. VRRP has the following characteristics and advantages:

- VRRP is a restartable application.
- When a VRRP master fails, the VRRP backup takes over within three times the advertisement time.
- VRRP over Ethernet, VRRP over VSAN, and Fibre Channel functions are implemented as defined in RFC 2338 and draft-ietf-vrrp-ipv6 spec, "Virtual Router Redundancy Propocol fol IPv6"
- A virtual router is mapped to each VSAN and Ethernet interface with its unique virtual router IP, virtual router MAC, and VR ID.
- Interface Mgmt 0 supports only one VRRP group. All other interface supports up to 7 virtual router groups, including both IPv4 and v6 combined.
- VR IDs can be reused in multiple VSANs with different virtual router IP mapping.
- Up to 255 virtual router groups can be assigned in each VSAN.
- VRRP security provides three options, including no authentication, simple text authentication, and MD5 authentication.
- Supports IPv4 and IPv6.

Note If you are using IPv6, you must either configure an IPv6 address on the interface or enable IPv6 on the interface. For more information about IPv6, see Chapter 48, "Configuring IPv6 for Gigabit Ethernet Interfaces."

In Figure 46-5, switch A is the VRRP master and switch B is the VRRP backup switch. Both switches have an IP address to VRRP mapping configured. The other switches set switch A as the default gateway. If switch A fails, the other switches do not have to change the routing configurations as switch B automatically becomes the master and takes over the function of a gateway.

Figure 46-5 VRRP Functionality

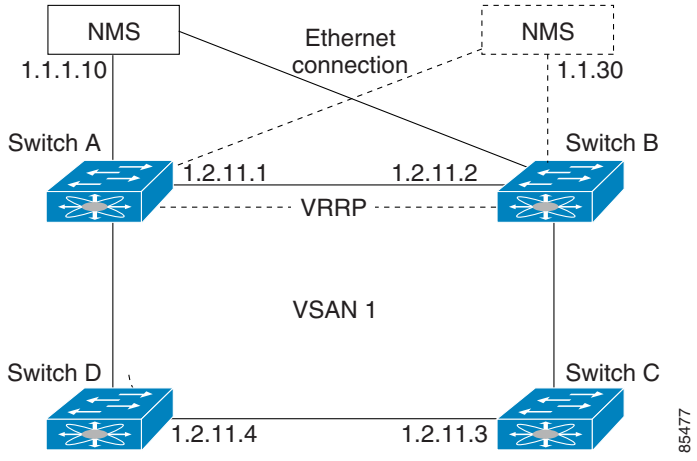

In Figure 46-6, the fabric example has two virtual router groups (VR1 and VR 2) because a virtual router cannot span across different types of interfaces. In both switch 1 and switch 2, the Ethernet interface is in VR 1 and the FC interface is in VR 2. Each virtual router is uniquely identified by the VSAN interface and the VR ID.

Figure 46-6 Redundant Gateway

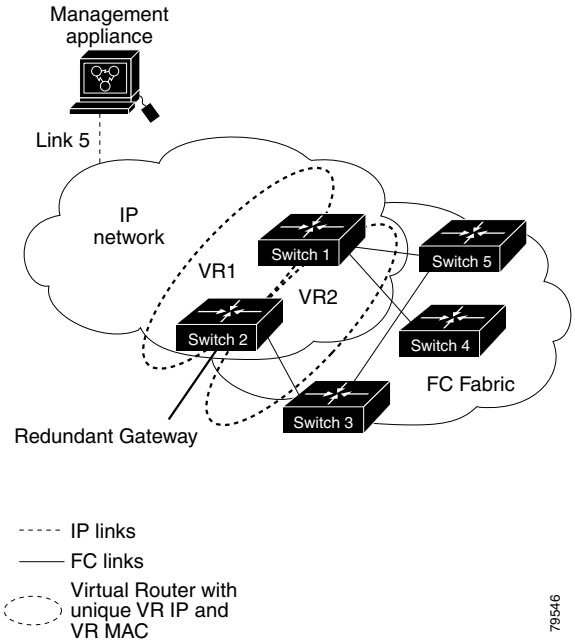

Configuring VRRP

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

Adding and Deleting Virtual Router

All VRRP configurations should be replicated across switches in a fabric that runs VRRP.

Note The total number of VRRP groups that you can configure on a Gigabit Ethernet port, including main interfaces and subinterfaces, cannot exceed seven. This limitation applies to both IPv4 and IPv6 groups.

Virtual Router Initiation

By default, a virtual router is always disabled. VRRP can be configured only if this state is enabled. Be sure to configure at least one IP address, either IPv4 or IPv6, before attempting to enable a VR.

Adding Virtual Router IP Addresses

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

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

To manage IP addresses for virtual routers from Device Manager, follow these steps:

Step 1 Choose **IP > VRRP**. You see the **Operations** tab of the VRRP dialog box.

Step 2 Click the **IP Addresses** tab on the VRRP dialog box.

Step 3 To create a new VRRP entry, click **Create**. You see the Create VRRP IP Addresses window.

Step 4 Complete the fields in this window to create a new VRRP IP Address, and click **OK** or **Apply**.

Priority for the Virtual Router

The valid range to assign a virtual router priority is 1 to 254 with 1 being the lowest priority and 254 being the highest priority. The default value is 100 for switches with secondary IP addresses and 255 for switches with the primary IP address.

Time Interval for Advertisement Packets

The valid time range for an advertisement packet on an interface using IPv4 is between 1 and 255 seconds. The default value is 1 (one) second. If the switch has the primary IP address, this time must be specified.

The valid time range for an advertisement packet on an interface using IPv6 is between 100 and 4095 centiseconds. The default value is 100 centisecond. If the switch has the primary IPF address, this time must be specified.

Priority Preemption

You can enable a higher priority backup virtual router to preempt the lower priority master virtual router.

Note If the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.

Note The VRRP preemption is not supported on IP storage Gigabit Ethernet interfaces.

Virtual Router Authentication

VRRP security provides three options, including simple text authentication, MD5 authentication, and no authentication.

- Simple text authentication uses a unique, 1 to 8 character password that is used by all switches participating in the same virtual router. This password should be different from other security passwords.
- MD5 authentication uses a unique, 16 character key that is shared by all switches participating in the same virtual router. This secret key is shared by all switches in the same virtual router.
- No authentication is the default option.

You can configure the key using the authentication option in the VRRP submode and distribute it using the configuration file. The security parameter index (SPI) settings assigned in this option should be unique for each VSAN.

Note All VRRP configurations must be duplicated.

Note VRRP router authentication does not apply to IPv6.

Priority Based on Interface State Tracking

The tracking feature is disabled by default. When you specify the tracking option, the priority of the virtual router is changed based on the state of another interface in the switch. When the tracked interface is down, the priority of the virtual router is changed to a lower priority value. When the tracked interface is up, the priority of the virtual router is restored to its original value. You can track one of two interfaces on a switch in the Cisco MDS 9000 Family: a specified VSAN interface or a management interface.

Note For interface tracking to function, you must enable preemption on the interface. See the "Priority Preemption" section on page 46-11.

DNS Server Configuration

The DNS client on the switch communicates with the DNS server to perform the IP address-name server correspondence.

The DNS server may be dropped after two attempts because of one of the following reasons:

- The IP address or the switch name is wrongly configured.
- The DNS server is not reachable because external reasons (reasons beyond our control).

Note When accessing a Telnet host, if the DNS server is not reachable (for any reason) the switch login prompt may take a longer time to appear. If so, verify that the DNS server is accurately configured and reachable.

Default Settings

Table 46-1 lists the default settings for VRRP features.

Table 46-1 Default VRRP Settings

Parameters	Default
Virtual router state	Disabled.
Maximum groups per VSAN	255.
Maximum groups per Gigabit Ethernet port	7.
Priority preemption	Disabled.
Virtual router priority	100 for switch with secondary IP addresses. 255 for switches with the primary IP address.
Priority interface tracking	Disabled.
Advertisement interval	1 second for IPv4. 100 centiseconds for IPv6.

Table 46-2 lists the default settings for DNS features.

Table 46-2 Default DNS Settings

Parameters	Default
Domain lookup	Disabled.
Domain name	Disabled.
Domains	None.
Domain server	None.
Maximum domain servers	6.

CHAPTER 47

Configuring IP Storage

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

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

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

This chapter includes the following sections:

Services Modules

The IP Storage services module (IPS module) and the MPS-14/2 module allow you to use FCIP and iSCSI features. Both modules integrate seamlessly into the Cisco MDS 9000 Family, and support the full range of features available on other switching modules, including VSANs, security, and traffic management. The following types of storage services modules are currently available for use in any switch in the Cisco MDS 9200 Series or in the Cisco MDS 9500 Series:

- The 4-port, hot-swappable IPS module (IPS-4) has four Gigabit Ethernet ports.
- The 8-port, hot-swappable IPS module (IPS-8) has eight Gigabit Ethernet ports.
- The MPS-14/2 module has 14 Fibre Channel ports (numbered 1 through 14) and two Gigabit Ethernet ports (numbered 1 and 2).

Gigabit Ethernet ports in these modules can be configured to support the FCIP protocol, the iSCSI protocol, or both protocols simultaneously.

- FCIP—FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices. Figure 47-1 shows how the IPS module is used in different FCIP scenarios.

Figure 47-1 FCIP Scenarios

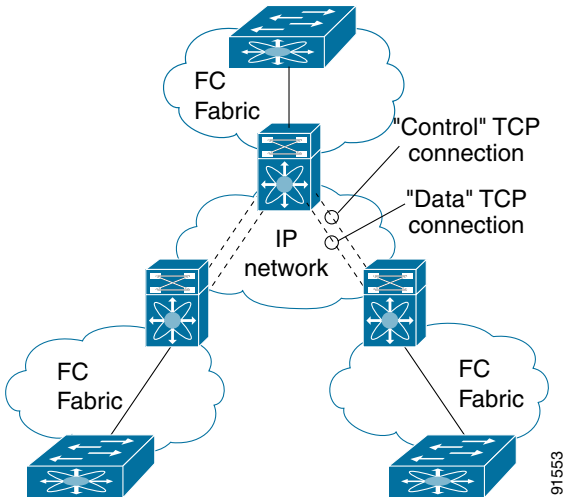

- iSCSI—The IPS module provides IP hosts access to Fibre Channel storage devices. The IP host sends SCSI commands encapsulated in iSCSI protocol data units (PDUs) to a Cisco MDS 9000 Family switch IPS port over a TCP/IP connection. At this point, the commands are routed from an IP network into a Fibre Channel network and forwarded to the intended target. Figure 47-2 depicts the iSCSI scenarios in which the IPS module is used.

Figure 47-2 iSCSI Scenarios

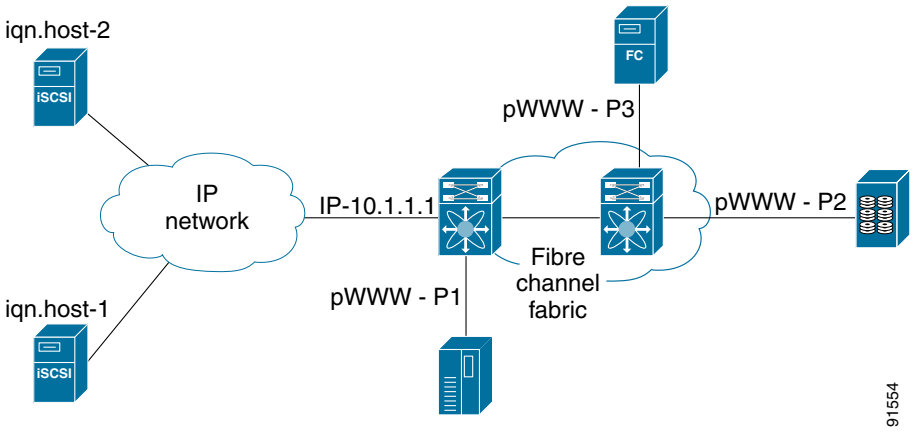

Module Status Verification

To verify the status of the module using Fabric Manager, follow these steps:

Step 1 Select a switch in the Fabric pane.

Step 2 Open the **Switches** folder and select **Hardware** in the Physical Attributes pane.

You see the status for all modules in the switch in the Information pane.

IPS Module Upgrade

Caution A software upgrade is only disruptive for the IPS module. The SAN-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

IPS modules use a rolling upgrade install mechanism where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each IPS module in a switch requires a 5-minute delay before the next IPS module is upgraded.

MPS-14/2 Module Upgrade

Caution A software upgrade is only partially disruptive for the MPS-14/2 module. The SAN-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

The MPS-14/2 modules have 14 Fibre Channel ports (nondisruptive upgrade) and 2 Gigabit Ethernet ports (disruptive upgrade). MPS-14/2 modules use a rolling upgrade install mechanism for the two Gigabit Ethernet ports where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each MPS-14/2 module in a switch requires a 5-minute delay before the next module is upgraded.

Supported Hardware

You can configure the FCIP and iSCSI features using one or more of the following hardware solutions:

- IPS-4 and IPS-8 modules (refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for more information)
- MPS-14/2 module (refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for more information).

Note In both the MPS-14/2 module and the Cisco MDS 9216i integrated supervisor module, the port numbering differs for the Fibre Channel ports and the Gigabit Ethernet ports. The Fibre Channel ports are numbered from 1 through 14 and the Gigabit Ethernet ports are numbered 1 and 2.

- Cisco MDS 9216i Switch (refer to the *Cisco MDS 9200 Series Hardware Installation Guide*).

Configuring Gigabit Ethernet Interfaces for IPv4

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each IPS module or MPS-14/2 module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured. This section covers the steps required to configure IP for subsequent use by FCIP and iSCSI.

Note For information about configuring FCIP, see Chapter 4, "Fabric Manager Client". For information about configuring iSCSI, see Chapter 45, "Configuring iSCSI".

A new port mode, called IPS, is defined for Gigabit Ethernet ports on each IPS module or MPS-14/2 module. IP storage ports are implicitly set to IPS mode, so it can only be used to perform iSCSI and FCIP storage functions. IP storage ports do not bridge Ethernet frames or route other IP packets.

Each IPS port represents a single virtual Fibre Channel host in the Fibre Channel SAN. All the iSCSI hosts connected to this IPS port are merged and multiplexed through the single Fibre Channel host.

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

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

Note To configure IPv6 on a Gigabit Ethernet interface, see the "Gigabit Ethernet IPv6-ACL Guidelines" section on page 48-23.

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

Basic Gigabit Ethernet Configuration

Figure 47-3 shows an example of a basic Gigabit Ethernet IP version 4 (IPv4) configuration.

Figure 47-3 Gigabit Ethernet IPv4 Configuration Example

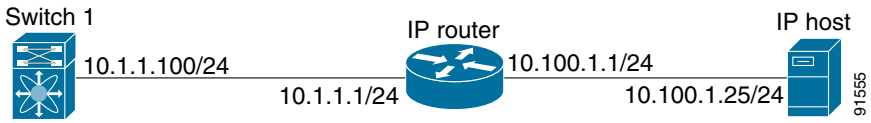

Note The port on the Ethernet switch to which the MDS Gigabit Ethernet interface is connected should be configured as a host port (also known as access port) instead of a switch port. Spanning tree configuration for that port (on the ethernet switch) should disabled. This helps avoid the delay in the management port coming up due to delay from Ethernet spanning tree processing that the Ethernet switch would run if enabled. FoR Cisco Ethernet switches, use either the **switchport host** command in IOS is or the **set port host** in Catalyst OS. Refer to the configuration guide for your Ethernet switch.

To configure the Gigabit Ethernet interface for the scenario in Figure 47-3, follow these steps

Step 1 From Fabric Manager, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane. You see the Gigabit Ethernet configuration in the Information pane.

From Device Manager, right-click the Gigabit Ethernet port that you want to configure and choose **Configure...**. You see the Gigabit Ethernet configuration dialog box.

Step 2 Click the **General** tab in Fabric Manager, or click the **GigE** tab in Device Manager to display the general configuration options for the interface.

Step 3 Set the description and MTU value for the interface. The valid value for the MTU field can be a number in the range from 576 to 9000.

Step 4 Set **Admin** up or down and check the **CDP** check box if you want this interface to participate in CDP.

Step 5 Set **IpAddress/Mask** with the IP address and subnet mask for this interface.

Step 6 From Fabric Manager, click the **Apply Changes** icon to save these changes, or click the **Undo Changes** icon to discard changes.

From Device Manager, click **Apply** to save these changes, or click **Close** to discard changes and close the Gigabit Ethernet configuration dialog box.

Configuring Interface Descriptions

See the "About Interface Modes" section on page 18-3 for details on configuring the switch port description for any interface.

Configuring Configuring Autonegotiation

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

Configuring the MTU Frame Size

You can configure the interfaces on a switch to transfer large (or jumbo) frames on a port. The default IP maximum transmission unit (MTU) frame size is 1500 bytes for all Ethernet ports. By configuring jumbo frames on a port, the MTU size can be increased up to 9000 bytes.

Note The minimum MTU size is 576 bytes.

Tip MTU changes are disruptive, all FCIP links and iSCSI sessions flap when the software detects a change in the MTU size.

Configuring Promiscuous Mode

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

About VLANs for Gigabit Ethernet

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

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

Note If the IPS module or MPS-14/2 module is connected to a Cisco Ethernet switch, and you need to have traffic from multiple VLANs coming to one IPS port, verify the following requirements on the Ethernet switch:
- The Ethernet switch port connected to the IPS module or MPS-14/2 module is configured as a trunking port.
- The encapsulation is set to 802.1Q and not ISL, which is the default.

Use the VLAN ID as a subscription to the Gigabit Ethernet interface name to create the subinterface name (the `<slot-number>/<port-number>.<VLAN-ID>`).

Interface Subnet Requirements

Gigabit Ethernet interfaces (major), subinterfaces (VLAN ID), and management interfaces (mgmt 0) can be configured in the same or different subnet depending on the configuration (see Table 47-1).

Table 47-1 Subnet Requirements for Interfaces

Interface 1	Interface 2	Same Subnet Allowed	Notes
Gigabit Ethernet 1/1	Gigabit Ethernet 1/2	Yes	Two major interfaces can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.100	Yes	Two subinterfaces with the same VLAN ID can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.200	No	Two subinterfaces with different VLAN IDs cannot be configured in the same subnet.
Gigabit Ethernet 1/1	Gigabit Ethernet 1/1.100	No	A subinterface cannot be configured on the same subnet as the major interface.
mgmt0	Gigabit Ethernet 1/1.100	No	The mgmt0 interface cannot be configured in the same subnet as the Gigabit Ethernet interfaces or subinterfaces.
mgmt0	Gigabit Ethernet 1/1	No	

Note The configuration requirements in Table 47-1 also apply to Ethernet PortChannels.

Verifying Gigabit Ethernet Connectivity

Once the Gigabit Ethernet interfaces are connected with valid IP addresses, verify the interface connectivity on each switch. Ping the IP host using the IP address of the host to verify that the static IP route is configured correctly.

Note If the connection fails, verify the following, and ping the IP host again:
- The IP address for the destination (IP host) is correctly configured.
- The host is active (powered on).
- The IP route is configured correctly.
- The IP host has a route to get to the Gigabit Ethernet interface subnet.
- The Gigabit Ethernet interface is in the `up` state.

Gigabit Ethernet IPv4-ACL Guidelines

Tip If IPv4-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group.

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

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

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

- Apply IPv4-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
  - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
  - The **established**, **precedence**, and **fragments** options are ignored when you apply IPv4-ACLs (containing these options) to Gigabit Ethernet interfaces.
  - If an IPv4-ACL rule applies to a preexisting TCP connection, that rule is ignored. For example if there is an existing TCP connection between A and B, and an IPv4-ACL specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.

Configuring Gigabit Ethernet High Availability

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

VRRP for iSCSI and FCIP Services

VRRP provides a redundant alternate path to the Gigabit Ethernet port for iSCSI and FCIP services. VRRP provides IP address failover protection to an alternate Gigabit Ethernet interface so the IP address is always available (see Figure 47-4).

Figure 47-4 VRRP Scenario

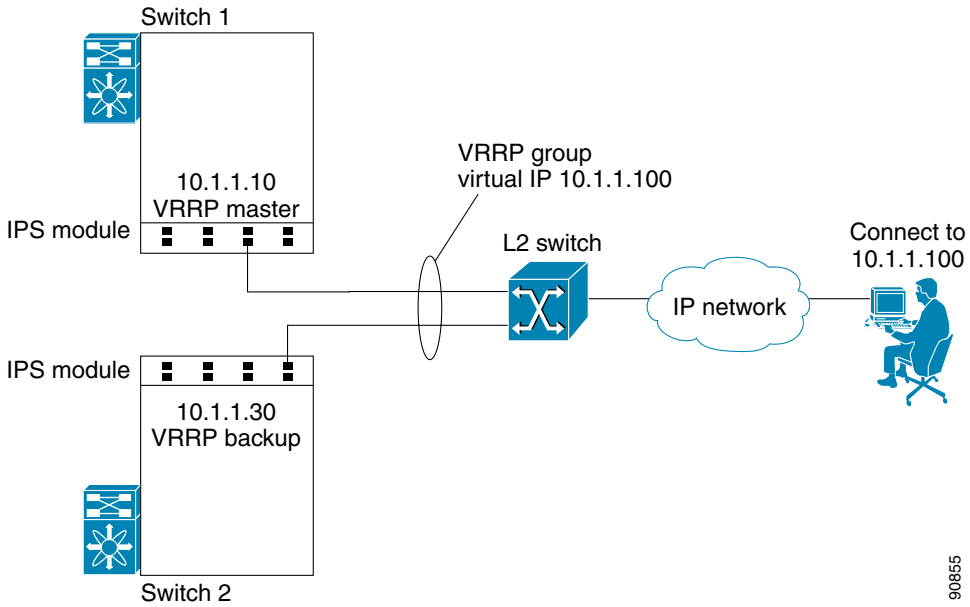

In Figure 47-4, all members of the VRRP group must be IP storage Gigabit Ethernet ports. VRRP group members can be one or more of the following interfaces:

- One or more interfaces in the same IPS module or MPS-14/2 module
- Interfaces across IPS modules or MPS-14/2 modules in one switch
- Interfaces across IPS modules or MPS-14/2 modules in different switches
- Gigabit Ethernet subinterfaces
- Ethernet PortChannels and PortChannel subinterfaces

See the "Virtual Router Redundancy Protocol" section on page 46-8.

Configuring VRRP for Gigabit Ethernet Interfaces

Note The VRRP **preempt** option is not supported on IPS Gigabit Ethernet interfaces. However, if the virtual IPv4 address is also the IPv4 address for the interface, then preemption is implicitly applied.

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

About Ethernet PortChannel Aggregation

Ethernet PortChannels refer to the aggregation of multiple physical Gigabit Ethernet interfaces into one logical Ethernet interface to provide link redundancy and, in some cases, higher aggregated bandwidth and load balancing.

An Ethernet switch connecting to the MDS switch Gigabit Ethernet port can implement load balancing based on the IP address, IP address and UDP/TCP port number, or MAC address. Due to the load balancing scheme, the data traffic from one TCP connection is always sent out on the same physical Gigabit Ethernet port of an Ethernet PortChannel. For the traffic coming to the MDS, an ethernet switch can implement load balancing based on its IP address, its source-destination MAC address, or its IP address and port. The data traffic from one TCP connection always travels on the same physical links. To make use of both ports for the outgoing direction, multiple TCP connections are required.

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

Note The Cisco Ethernet switch's PortChannel should be configured as a static PortChannel, and not the default 802.3ad protocol.

Ethernet PortChannels can only aggregate two physical interfaces that are adjacent to each other on a given IPS module (see Figure 47-5).

Note PortChannel members must be one of these combinations: ports 1–2, ports 3–4, ports 5–6, or ports 7–8.

Figure 47-5 Ethernet PortChannel Scenario

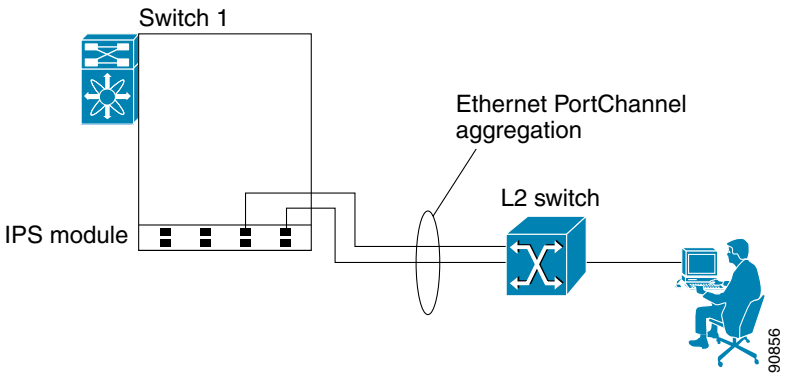

In Figure 47-5, Gigabit Ethernet ports 3 and 4 in slot 9 are aggregated into an Ethernet PortChannel. Ethernet PortChannels are not supported on MPS-14/2 modules and 9216i IPS modules.

Note PortChannel interfaces provide configuration options for both Gigabit Ethernet and Fibre Channel. However, based on the PortChannel membership, only Gigabit Ethernet parameters or Fibre Channel parameters are applicable.

Configuring Ethernet PortChannels

The PortChannel configuration specified in Chapter 21, "Configuring PortChannels" also applies to Ethernet PortChannel configurations.

Note Gigabit Ethernet interfaces cannot be added to a PortChannel if one of the following cases apply:
- The interface already has an IP address assigned.
- The subinterfaces are configured on that interface.
- The interface already has an associated IPv4-ACL rule and the PortChannel does not.

Configuring CDP

The Cisco Discovery Protocol (CDP) is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interfaces on the IPS module or MPS-14/2 module.

See the "Configuring CDP" section on page 11-11.

CHAPTER 48

Configuring IPv6 for Gigabit Ethernet Interfaces

IP version 6 (IPv6) provides extended addressing capability beyond those provided in IP version 4 (IPv4) in Cisco MDS SAN-OS. The architecture of IPv6 has been designed to allow existing IPv4 users to transition easily to IPv6 while providing services such as end-to-end security, quality of service (QoS), and globally unique addresses.

This chapter includes the following sections:

About IPv6

IPv6 provides the following enhancements over IPv4:

- Allows networks to scale and provide global reachability.
- Reduces the need for private address and network address translation (NAT).
- Provides simpler autoconfiguration of addresses.

This section describes the IPv6 features supported by Cisco MDS SAN-OS and includes the following topics:

Extended IPv6 Address Space for Unique Addresses

IPv6 extends the address space by quadrupling the number of network address bits from 32 bits (in IPv4) to 128 bits, which provides many more globally unique IP addresses. By being globally unique, IPv6 addresses enable global reachability and end-to-end security for networked devices, functionality that is crucial to the applications and services that are driving the demand for more addresses.

IPv6 Address Formats

IPv6 addresses are represented as a series of 16-bit hexadecimal fields separated by colons (:) in the format x:x:x:x:x:x:x:x. The following are examples of IPv6 addresses:

2001:0DB8:7654:3210:FEDC:BA98:7654:3210

2001:0DB8:0:0:8:800:200C:417A

It is common for IPv6 addresses to contain successive hexadecimal fields of zeros. To make IPv6 addresses easier to use, two colons (::) may be used to compress successive hexadecimal fields of zeros at the beginning, middle, or end of an IPv6 address (the colons represent successive hexadecimal fields of zeros). Table 48-1 lists compressed IPv6 address formats.

Note Two colons (::) can be used only once in an IPv6 address to represent the longest successive hexadecimal fields of zeros.

Note The hexadecimal letters in IPv6 addresses are not case-sensitive.

Table 48-1 Compressed IPv6 Address Formats

IPv6 Address Type	Uncompressed Format	Compressed Format
Unicast	2001:0DB8:800:200C:0:0:0:417A	2001:0DB8:800:200C::417A
Multicast	FF01:0:0:0:0:0:0:101	FF01::101

IPv6 Address Prefix Format

An IPv6 address prefix, in the format *ipv6-prefix/prefix-length*, can be used to represent bit-wise contiguous blocks of the entire address space. The *ipv6-prefix* is specified in hexadecimal using 16-bit values between the colons. The *prefix-length* is a decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). For example, 2001:0DB8:8086:6502::/32 is a valid IPv6 prefix.

IPv6 Address Type: Unicast

An IPv6 unicast address is an identifier for a single interface on a single node. A packet that is sent to a unicast address is delivered to the interface identified by that address. The Cisco MDS SAN-OS supports the following IPv6 unicast address types:

- Global addresses
- Link-local addresses

Global Addresses

Global IPv6 addresses are defined by a global routing prefix, a subnet ID, and an interface ID. Figure 48-1 shows the structure of a global address.

Figure 48-1 Global Address Format

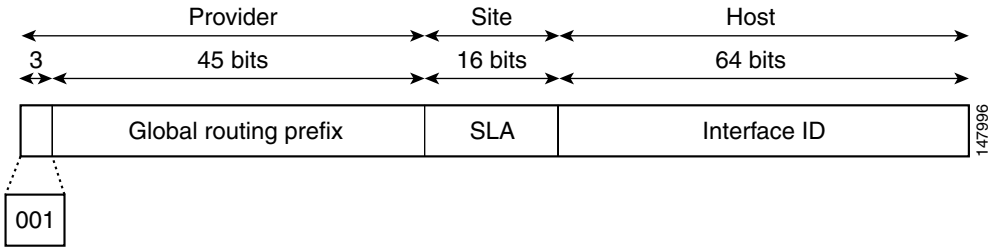

Addresses with a prefix of 2000::/3 (001) through E000::/3 (111) are required to have 64-bit interface identifiers in the extended universal identifier (EUI)-64 format. The Internet Assigned Numbers Authority (IANA) allocates the IPv6 address space in the range of 2000::/16 to regional registries.

The aggregatable global address typically consists of a 48-bit global routing prefix and a 16-bit subnet ID or Site-Level Aggregator (SLA). In the IPv6 aggregatable global unicast address format document (RFC 2374), the global routing prefix included two other hierarchically structured fields named Top-Level Aggregator (TLA) and Next-Level Aggregator (NLA).The IETF decided to remove the TLS and NLA fields from the RFCs because these fields are policy-based. Some existing IPv6 networks deployed before the change might still be using networks based on the older architecture.

A 16-bit subnet field called the subnet ID could be used by individual organizations to create their own local addressing hierarchy and to identify subnets. A subnet ID is similar to a subnet in IPv4, except that an organization with an IPv6 subnet ID can support up to 65,535 individual subnets.

An interface ID is used to identify interfaces on a link. The interface ID must be unique to the link. They may also be unique over a broader scope. In many cases, an interface ID will be the same as, or based on, the link-layer address of an interface, which results in a globally unique interface ID. Interface IDs used in aggregatable global unicast and other IPv6 address types must be 64 bits long and constructed in the modified EUI-64 format.

Cisco MDS SAN-OS supports IEEE 802 interface types (for example, Gigabit Ethernet interfaces). The first three octets (24 bits) are taken from the Organizationally Unique Identifier (OUI) of the 48-bit link-layer address (MAC address) of the interface, the fourth and fifth octets (16 bits) are a fixed hexadecimal value of FFFE, and the last three octets (24 bits) are taken from the last three octets of the MAC address. The construction of the interface ID is completed by setting the Universal/Local (U/L) bit—the seventh bit of the first octet—to a value of 0 or 1. A value of 0 indicates a locally administered identifier; a value of 1 indicates a globally unique IPv6 interface identifier (see Figure 48-2).

Figure 48-2 Interface Identifier Format

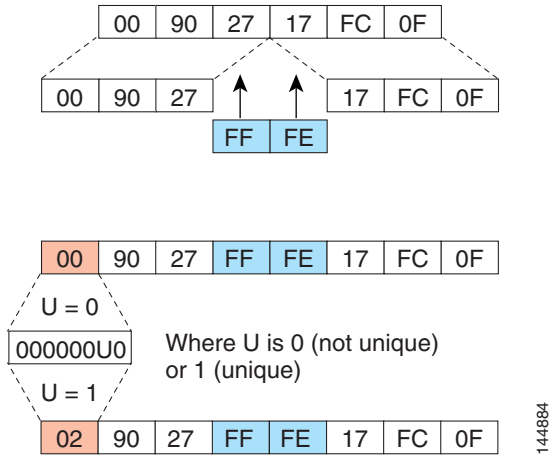

Link-Local Address

A link-local address is an IPv6 unicast address that is automatically configured on an interface using the link-local prefix FE80::/10 and the interface identifier in the modified EUI-64 format. Link-local addresses are used in the neighbor discovery protocol and the stateless autoconfiguration process. Nodes on a local link can use link-local addresses to communicate. Figure 48-3 shows the structure of a link-local address.

Figure 48-3 Link-Local Address Format

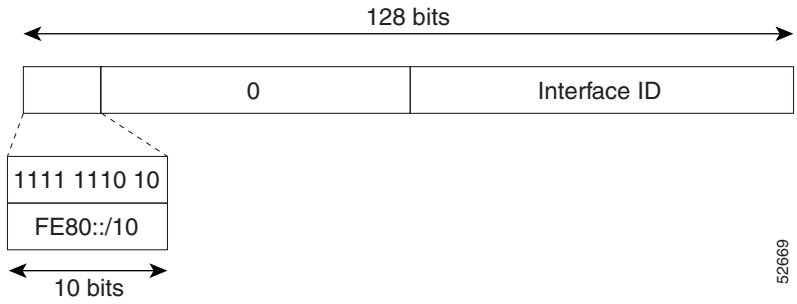

IPv6 Address Type: Multicast

An IPv6 multicast address is an IPv6 address that has a prefix of FF00::/8 (1111 1111). An IPv6 multicast address is an identifier for a set of interfaces that typically belong to different nodes. A packet sent to a multicast address is delivered to all interfaces identified by the multicast address. The second octet following the prefix defines the lifetime and scope of the multicast address. A permanent multicast address has a lifetime parameter equal to 0; a temporary multicast address has a lifetime parameter equal to 1. A multicast address has the scope of a node, link, site, or organization, or a global scope has a scope parameter of 1, 2, 5, 8, or E, respectively. For example, a multicast address with the prefix FF02::/16 is a permanent multicast address with a link scope. Figure 48-4 shows the format of the IPv6 multicast address.

Figure 48-4 IPv6 Multicast Address Format

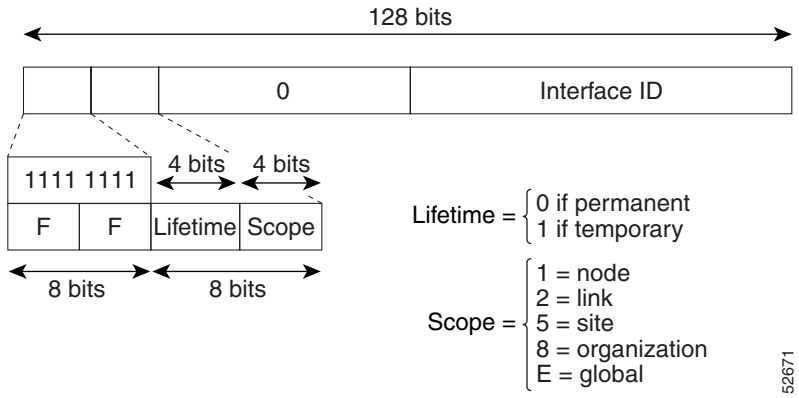

IPv6 hosts are required to join (receive packets destined for) the following multicast groups:

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

  The solicited-node multicast address is a multicast group that corresponds to an IPv6 unicast address. IPv6 nodes must join the associated solicited-node multicast group for every unicast address to which it is assigned. The IPv6 solicited-node multicast address has the prefix FF02:0:0:0:0:1:FF00:0000/104 concatenated with the 24 low-order bits of a corresponding IPv6 unicast address. (See Figure 48-5.) For example, the solicited-node multicast address corresponding to the IPv6 address 2037::01:800:200E:8C6C is FF02::1:FF0E:8C6C. Solicited-node addresses are used in neighbor solicitation messages.

Figure 48-5 IPv6 Solicited-Node Multicast Address Format

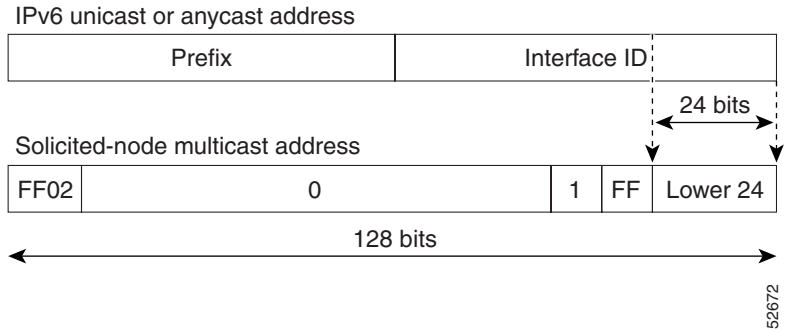

Note There are no broadcast addresses in IPv6. IPv6 multicast addresses are used instead of broadcast addresses.

ICMP for IPv6

Internet Control Message Protocol (ICMP) in IPv6 functions the same as ICMP in IPv4—ICMP generates error messages such as ICMP destination unreachable messages, and informational messages such as ICMP echo request and reply messages. Additionally, ICMP packets in IPv6 are used in the IPv6 neighbor discovery process, path MTU discovery, and the Multicast Listener Discovery (MLD) protocol for IPv6. MLD is based on version 2 of the Internet Group Management Protocol (IGMP) for IPv4.

A value of 58 in the Next Header field of the basic IPv6 packet header identifies an IPv6 ICMP packet. ICMP packets in IPv6 resemble a transport-layer packet in the sense that the ICMP packet follows all the extension headers and is the last piece of information in the IPv6 packet. Within IPv6 ICMP packets, the ICMPv6 Type and ICMPv6 Code fields identify IPv6 ICMP packet specifics, such as the ICMP message type. The value in the Checksum field is derived (computed by the sender and checked by the receiver) from the fields in the IPv6 ICMP packet and the IPv6 pseudoheader. The ICMPv6 Data field contains error or diagnostic information relevant to IP packet processing. Figure 48-6 shows the IPv6 ICMP packet header format.

Figure 48-6 IPv6 ICMP Packet Header Format

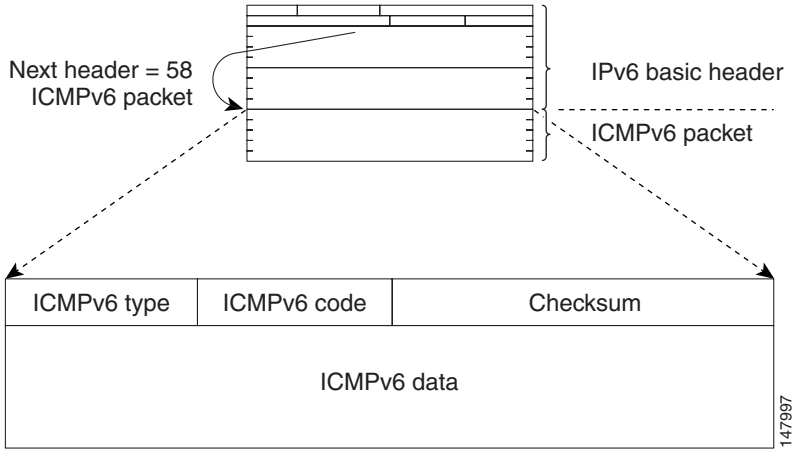

Path MTU Discovery for IPv6

As in IPv4, path MTU discovery in IPv6 allows a host to dynamically discover and adjust to differences in the MTU size of every link along a given data path. In IPv6, however, fragmentation is handled by the source of a packet when the path MTU of one link along a given data path is not large enough to accommodate the size of the packets. Having IPv6 hosts handle packet fragmentation saves IPv6 router processing resources and helps IPv6 networks run more efficiently.

Note In IPv4, the minimum link MTU is 68 octets, which means that the MTU size of every link along a given data path must support an MTU size of at least 68 octets.

In IPv6, the minimum link MTU is 1280 octets. We recommend using an maximum transmission unit (MTU) value of 1500 octets for IPv6 links.

IPv6 Neighbor Discovery

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

IPv6 Neighbor Solicitation and Advertisement Messages

A value of 135 in the Type field of the ICMP packet header identifies a neighbor solicitation message. Neighbor solicitation messages are sent on the local link when a node wants to determine the link-layer address of another node on the same local link. (See Figure 48-7.) When a node wants to determine the link-layer address of another node, the source address in a neighbor solicitation message is the IPv6 address of the node sending the neighbor solicitation message. The destination address in the neighbor solicitation message is the solicited-node multicast address that corresponds to the IPv6 address of the destination node. The neighbor solicitation message also includes the link-layer address of the source node.

Figure 48-7 IPv6 Neighbor Discovery—Neighbor Solicitation Message

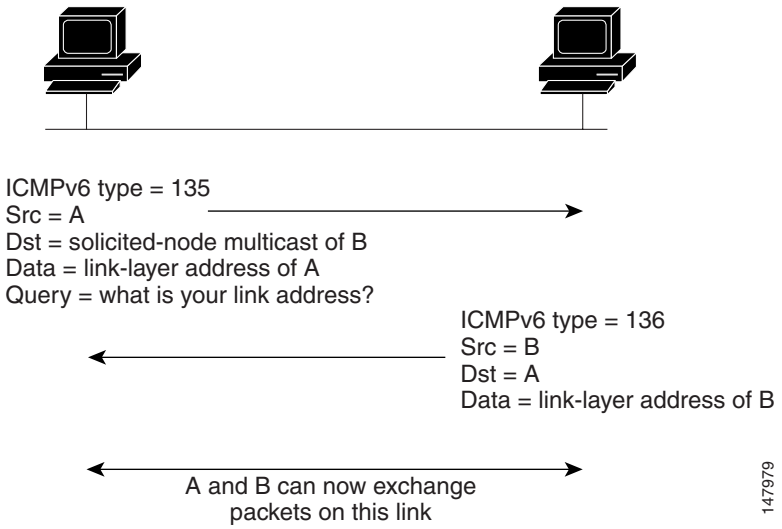

After receiving the neighbor solicitation message, the destination node replies by sending a neighbor advertisement message, which has a value of 136 in the Type field of the ICMP packet header, on the local link. The source address in the neighbor advertisement message is the IPv6 address of the node (more specifically, the IPv6 address of the node interface) sending the neighbor advertisement message. The destination address in the neighbor advertisement message is the IPv6 address of the node that sent the neighbor solicitation message. The data portion of the neighbor advertisement message includes the link-layer address of the node sending the neighbor advertisement message.

After the source node receives the neighbor advertisement, the source node and destination node can communicate.

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. When a node wants to verifying the reachability of a neighbor, the destination address in a neighbor solicitation message is the unicast address of the neighbor.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link. When there is such a change, the destination address for the neighbor advertisement is the all-node multicast address.

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. Neighbor unreachability detection identifies the failure of a neighbor or the failure of the forward path to the neighbor, and is used for all paths between hosts and neighboring nodes (hosts or routers). Neighbor unreachability detection is performed for neighbors to which only unicast packets are being sent and is not performed for neighbors to which multicast packets are being sent.

A neighbor is considered reachable when the neighbor returns a positive acknowledgment indicating that it has received and processed packets previously sent to it. A positive acknowledgment could be from an upper-layer protocol such as TCP indicating that a connection is making forward progress (reaching its destination) or the receipt of a neighbor advertisement message in response to a neighbor solicitation message. If packets are reaching the peer, they are also reaching the next-hop neighbor of the source. Therefore, forward progress is also a confirmation that the next-hop neighbor is reachable.

For destinations that are not on the local link, forward progress implies that the first-hop router is reachable. When acknowledgments from an upper-layer protocol are not available, a node probes the neighbor using unicast neighbor solicitation messages to verify that the forward path is still working. The return of a solicited neighbor advertisement message from the neighbor is a positive acknowledgment that the forward path is still working (neighbor advertisement messages that have the solicited flag set to a value of 1 are sent only in response to a neighbor solicitation message). Unsolicited messages confirm only the one-way path from the source to the destination node; solicited neighbor advertisement messages indicate that a path is working in both directions.

Note A neighbor advertisement message that has the solicited flag set to a value of 0 must not be considered as a positive acknowledgment that the forward path is still working.

Neighbor solicitation messages are also used in the stateless autoconfiguration process to verify the uniqueness of unicast IPv6 addresses before the addresses are assigned to an interface. Duplicate address detection is performed first on a new, link-local IPv6 address before the address is assigned to an interface (the new address remains in a tentative state while duplicate address detection is performed). Specifically, a node sends a neighbor solicitation message with an unspecified source address and a tentative link-local address in the body of the message. If another node is already using that address, the node returns a neighbor advertisement message that contains the tentative link-local address. If another node is simultaneously verifying the uniqueness of the same address, that node also returns a neighbor solicitation message. If no neighbor advertisement messages are received in response to the neighbor solicitation message and no neighbor solicitation messages are received from other nodes that are attempting to verify the same tentative address, the node that sent the original neighbor solicitation message considers the tentative link-local address to be unique and assigns the address to the interface.

Every IPv6 unicast address (global or link-local) must be checked for uniqueness on the link; however, until the uniqueness of the link-local address is verified, duplicate address detection is not performed on any other IPv6 addresses associated with the link-local address.

Router Discovery

Router discovery performs both router solicitation and router advertisement. Router solicitations are sent by hosts to all-routers multicast addresses. Router advertisements are sent by routers in response to solicitations or unsolicited and contain default router information as well as additional parameters such as the MTU and hop limit.

IPv6 Stateless Autoconfiguration

All interfaces on IPv6 nodes must have a link-local address, which is automatically configured from the identifier for an interface and the link-local prefix FE80::/10. A link-local address enables a node to communicate with other nodes on the link and can be used to further configure the node.

Nodes can connect to a network and automatically generate site-local and global IPv6 address without the need for manual configuration or help of a server, such as a DHCP server. With IPv6, a router on the link advertises in router advertisement (RA) messages any site-local and global prefixes, and its willingness to function as a default router for the link. RA messages are sent periodically and in response to router solicitation messages, which are sent by hosts at system startup. (See Figure 48-8.)

Figure 48-8 IPv6 Stateless Autoconfiguration

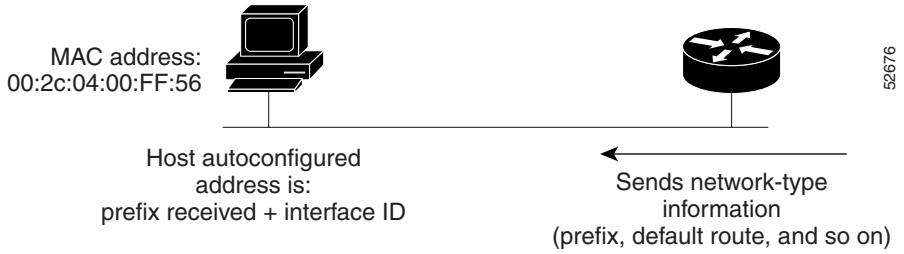

A node on the link can automatically configure site-local and global IPv6 addresses by appending its interface identifier (64 bits) to the prefixes (64 bits) included in the RA messages. The resulting 128-bit IPv6 addresses configured by the node are then subjected to duplicate address detection to ensure their uniqueness on the link. If the prefixes advertised in the RA messages are globally unique, then the IPv6 addresses configured by the node are also guaranteed to be globally unique. Router solicitation messages, which have a value of 133 in the Type field of the ICMP packet header, are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled RA message.

Dual IPv4 and IPv6 Protocol Stacks

The dual IPv4 and IPv6 protocol stack technique is one technique for a transition to IPv6. It enables gradual, one-by-one upgrades to applications running on nodes. Applications running on nodes are upgraded to make use of the IPv6 protocol stack. Applications that are not upgraded—they support only the IPv4 protocol stack—can coexist with upgraded applications on the same node. New and upgraded applications simply make use of both the IPv4 and IPv6 protocol stacks. (See Figure 48-9.)

Figure 48-9 Dual IPv4 and IPv6 Protocol Stack Technique

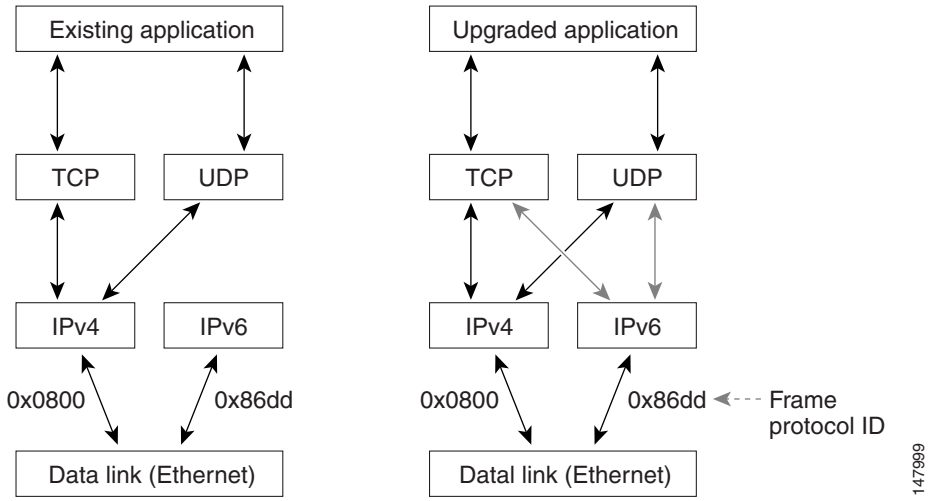

A new API has been defined to support both IPv4 and IPv6 addresses and DNS requests. An application can be upgraded to the new API and still use only the IPv4 protocol stack. The Cisco MDS SAN-OS supports the dual IPv4 and IPv6 protocol stack technique. When an interface is configured with both an IPv4 and an IPv6 address, the interface will accept and process both IPv4 and IPv6 traffic.

In Figure 48-10, an application that supports dual IPv4 and IPv6 protocol stacks requests all available addresses for the destination host name www.a.com from a DNS server. The DNS server replies with all available addresses (both IPv4 and IPv6 addresses) for www.a.com. The application chooses an address—in most cases, IPv6 addresses are the default choice—and connects the source node to the destination using the IPv6 protocol stack.

Figure 48-10 Dual IPv4 and IPv6 Protocol Stack Applications

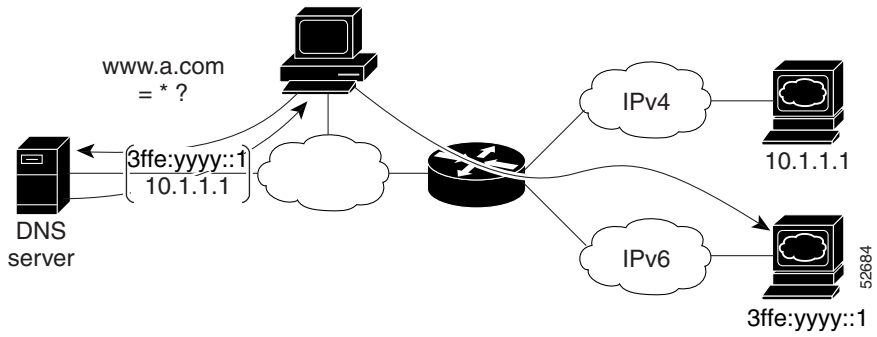

Configuring Basic Connectivity for IPv6

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

Configuring IPv6 Addressing and Enabling IPv6 Routing

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

You can configure IPv6 addresses on the following interface types:

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

Note The IPv6 address must be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons.

The ipv6-prefixmust be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons.

The IPv6 prefix-length is a decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). A slash mark must precede the decimal value.

Configuring a global IPv6 address on an interface automatically configures a link-local address and activates IPv6 for that interface. Additionally, the configured interface automatically joins the following required multicast groups for that link:

- Solicited-node multicast group FF02:0:0:0:0:1:FF00::/104 for each unicast address assigned to the interface
- All-node link-local multicast group FF02::1

Note The solicited-node multicast address is used in the neighbor discovery process.

Note The maximum number of IPv6 addresses (static and autoconfigured) allowed on an interface is eight, except on the management (mgmt 0) interface where only one static IPv6 address can be configured.

To configure an IPv6 address on an interface using Device Manager, follow these steps:

Step 1 Choose **Interfaces > Gigabit Ethernet and iSCSI**.

You see the Gigabit Ethernet Configuration dialog box (see Figure 48-11).

Figure 48-11 Gigabit Ethernet Configuration in Device Manager

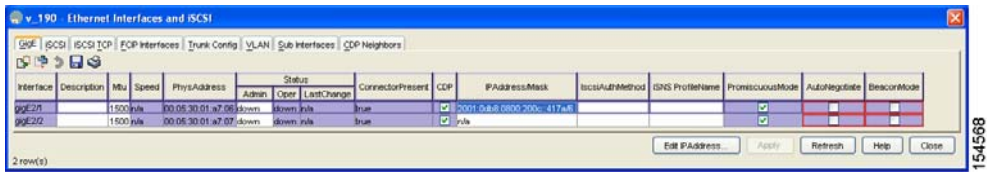

Step 2 Click the IP Address that you want to configure and click **Edit IP Address**.

You see the IP Address dialog box.

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

Step 4 Click **Create** to save these changes or click **Close** to discard any unsaved changes.

To enable IPv6 routing using Device Manager, follow these steps:

Step 1 Choose **IP > Routing**. You see the IP Routing Configuration dialog box(see Figure 48-11).

Figure 48-12 IP Routing Configuration in Device Manager

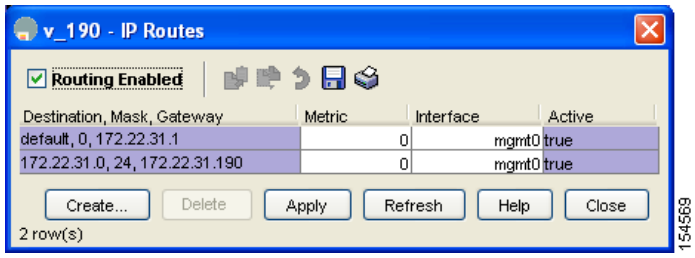

Step 2 Check the **Routing Enabled** check box.

Step 3 Click **Apply** to save these changes or click **Close** to discard any unsaved changes.

Configuring IPv4 and IPv6 Protocol Addresses

When an interface in a Cisco networking device is configured with both an IPv4 and an IPv6 address, the interface can send and receive data on both IPv4 and IPv6 networks.

To configure an interface in a Cisco networking device to support both the IPv4 and IPv6 protocol stacks using Device Manager, follow these steps:

Step 1 Click **Interfaces > Gigabit Ethernet and iSCSI**.

You see the Gigabit Ethernet Configuration dialog box.

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

You see the IP Address dialog box.

Step 3 Click **Create** and set the IP Address/Mask field, using the IPv4 or IPv6 format.

Step 4 Click **Create** to save these changes or click **Close** to discard any unsaved changes.

Configuring IPv6 Static Routes

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

Configuring a IPv6 Static Route

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

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

Step 1 Choose **IP > Routing**.

You see the IP Routing Configuration dialog box.

Step 2 Click **Create**.

You see the Create IP Route dialog box.

Step 3 Set the Dest field to the IPv6 destination address.

Step 4 Set the Mask field to the IPv6 subnet mask.

Step 5 Set the Gateway field to the IPv6 default gateway.

Step 6 Optionally, set the Metric field to the desired route metric.

Step 7 Select the interface from the Interface drop-down menu.

Step 8 Click **Create** to save these changes or click **Close** to discard any unsaved changes.

Gigabit Ethernet IPv6-ACL Guidelines

Tip If IPv6-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to a Ethernet PortChannel group. See Chapter 37, "Configuring IPv6 Access Control Lists," for information on configuring IPv6-ACLs.

Follow these guidelines when configuring IPv6-ACLs for Gigabit Ethernet interfaces:

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

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

- Apply IPv6-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
  - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
  - The **established**, **precedence**, and **fragments** options are ignored when you apply IPv6-ACLs (containing these options) to Gigabit Ethernet interfaces.
  - If an IPv6-ACL rule applies to a preexisting TCP connection, that rule is ignored. For example, if there is an existing TCP connection between A and B and an IPv6-ACL that specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.

See Chapter 37, "Configuring IPv6 Access Control Lists" for information on applying IPv6-ACLs to an interface.

Transitioning from IPv4 to IPv6

Cisco MDS SAN-OS does not support any transitioning mechanisms from IPv4 to IPv6. However, you can use the transitioning schemes in the Cisco router products for this purpose. For information on configuring Cisco routers to transition your network, refer to the *Implementing Tunneling for IPv6* document in the *Cisco IOS IPv6 Configuration Library*.

Default Settings

Table 48-2 lists the default settings for IPv6 parameters.

Table 48-2 Default IPv6 Parameters

Parameters	Default
IPv6 processing	Disabled.
Duplicate address detection attempts	0 (neighbor discovery disabled).
Reachability time	1000 milliseconds.
Retransmission time	30000 milliseconds.
IPv6-ACLs	None.

Cisco Systems

Send documentation comments to mdsfeedback-doc@cisco.com

PART 7

Intelligent Storage Services

CHAPTER 49

Configuring SCSI Flow Services and Statistics

Intelligent Storage Services are features supported on the Storage Services Module (SSM). Intelligent Storage Services supported in Cisco MDS SAN-OS Release 2.0(2b) and later include the following topics:

- SCSI flow services
- SCSI flow statistics

This chapter includes the following sections:

Enabling Intelligent Storage Services

In Cisco MDS SAN-OS Release 2.1(1a) or later, you can provision a subset of the ports for an SSM feature. The port range must be a multiple of four (for example fc4/1 through fc4-12).

To enable Intelligent Storage Services for an SSM and provision all ports or a group of ports to use these services using Fabric Manager, follow these steps:

Step 1 Expand **End Devices** and then select **SSM Features** in the Physical Attributes pane.

You see the Intelligent Storage Services configuration in the Information pane.

Step 2 Click the **SSM** tab.

You see the set of configured services in the Information pane shown in Figure 49-1.

Figure 49-1 SSM Configured Services

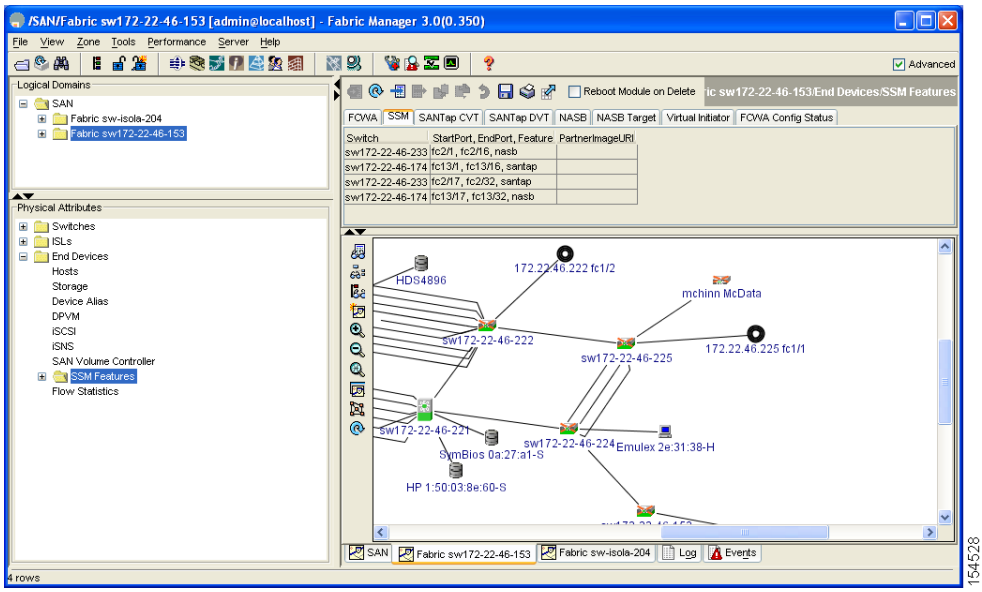

Step 3 Click **Create Row** to enable a new service on an SSM.

You see the Create SSM dialog box shown in Figure 49-2.

Figure 49-2 Create SSM Dialog Box

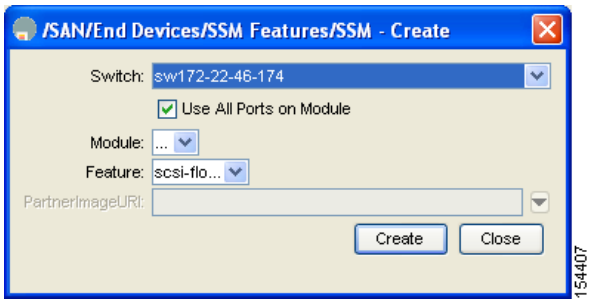

Step 4 Select the switch and SSM card you want to configure.

Step 5 Optionally, uncheck the **Use All Ports on Module** check box if you want to provision a subset of the ports on the card to use this service.

Step 6 Select the port range you want to provision for using this service (starting port and ending port).

Note The port range must be a multiple of four (for example fc4/1 through fc4-12).

Step 7 Select the feature you want to enable on these ports from the drop-down list of services.

Step 8 Set the PartnerImageURI field if you are enabling a third-party application that requires an image loaded onto the SSM.

Step 9 Click **Create** to create this row and enable this service or click **Cancel** to discard all changes.

Disabling Intelligent Storage Services

To disable Intelligent Storage Services in Fabric Manager for an SSM and free up a group of ports that used these services, follow these steps:

Step 1 Expand **End Devices** and then select **SSM Features** in the Physical Attributes pane.

You see the Intelligent Storage Services configuration in the Information pane.

Step 2 Click the **SSM** tab.

You see the set of configured services in the Information pane shown in Figure 49-3.

Figure 49-3 SSM Configured Services

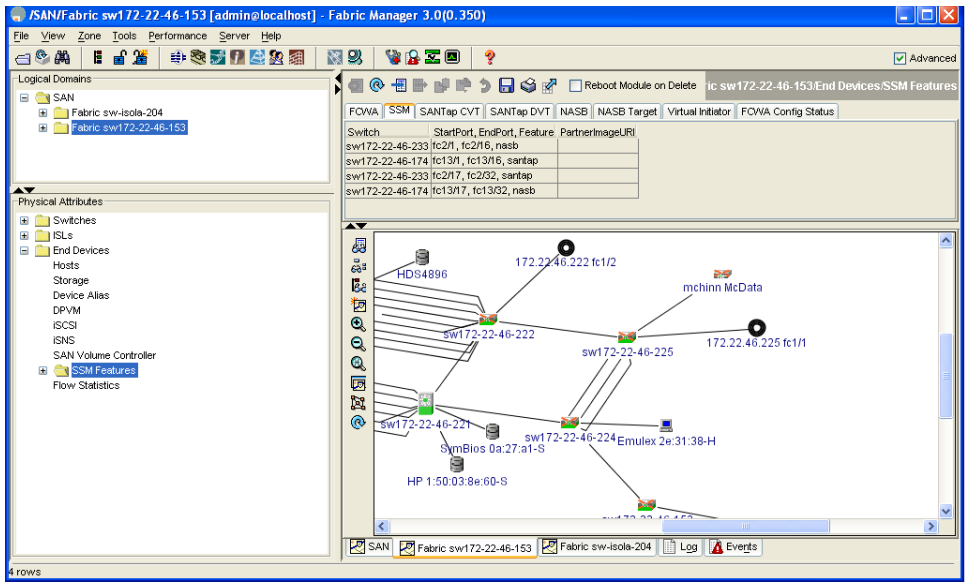

Step 3 Select the row in the table that you want to disable.

Step 4 Optionally, check the **Reboot Module on Delete** check box if you want to force the card to reboot after disabling the service (see Figure 49-3, next to the icons). This is equivalent to the CLI **force** option.

Step 5 Click **Delete Row**. The ports that were provisioned for this service become available for provisioning in another service.

Note If **Reboot Module on Delete** was checked, then the SSM module reboots.

About SCSI Flow Services

A SCSI initiator/target combination is a SCSI flow. SCSI flow services provide enhanced features for SCSI flows, such as write acceleration and flow monitoring for statistics gathering on an SSM.

Functionally, the SCSI flow services functional architecture consists of the following components:

- SCSI flow manager (SFM) on the supervisor
- SCSI flow configuration CLI on the supervisor
- SCSI flow configuration client on the Control Path Processor (CPP) of an SSM
- SCSI flow feature set support on the Data Path Processor (DPP) of an SSM

Figure 49-4 shows an example of the SCSI flow services functional architecture.

Figure 49-4 SCSI Flow Services Functional Architecture

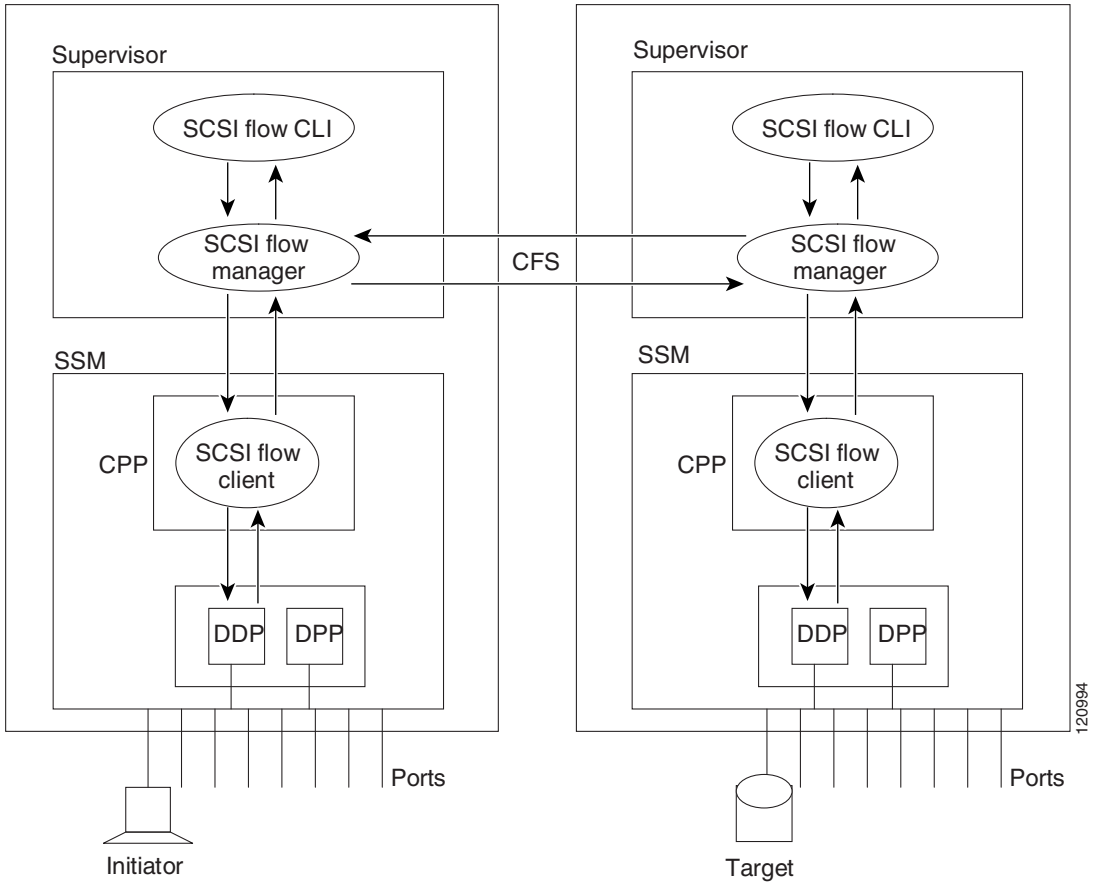

Note The SCSI target and initiator must be connected to different SSMs on different switches.

Note For statistics monitoring, the target device is not required to be connected to an SSM.

SCSI Flow Manager

The SCSI flow manager (SFM) resides on a supervisor module and handles the configuration of SCSI flows, validating them and relaying configuration information to the appropriate SSM. It also handles any dynamic changes to the status of the SCSI flow due to external events. The SFM registers events resulting from operations, such as port up or down, VSAN suspension, and zoning that affects the SCSI flow status, and updates the flow status and configuration accordingly.

The SFM on the initiator communicates to its peer on the target side using Cisco Fabric Services (CFS). Peer communication allows the initiator SFM to validate target parameters and program information on the target side.

SCSI Flow Configuration Client

A SCSI flow configuration client (SFCC) resides on the CPP of the SSM. It receives flow configuration requests from the SFM, programs the DPP corresponding to the initiator and target port interfaces, and responds to the SFM with the status of the configuration request.

SCSI Flow Data Path Support

The DPP on the SSM examines all the messages between the initiator and target and provides SCSI flow features such as Fibre Channel write acceleration and statistics monitoring.

Configuring SCSI Flow Services

A SCSI flow specification consists of the following attributes:

- SCSI flow identifier
- VSAN identifier
- SCSI initiator port WWN
- SCSI target port WWN
- Flow feature set consisting of Fibre Channel write acceleration and statistics monitoring.

The SCSI flow specification is a distributed configuration because the SCSI initiator and the target might be physically connected to SSMs on two different switches located across the fabric. The configuration does not require information to identify either the switch name or the SSM slot location for either the initiator or the target. The manual SCSI flow configuration is performed only at the initiator side. This simplifies the configuration process. The initiator switch sends the configuration to the SFM on the target switch using CFS. No SCSI flow configuration is necessary on the target switch.

Enabling SCSI Flow Services

In Cisco MDS SAN-OS Releases 2.0(1b) through 2.1(1a), you can only enable SCSI flow services on the entire SSM. As of Cisco MDS SAN-OS Release 2.1(2), you can enable SCSI flow services either on the entire SSM or on groups of four interfaces.

Enabling SCSI flow services on interfaces has the following restrictions:

- The fewest number of interfaces which you can enable is four. You can specify fc1 through fc4 but not fc1 through fc2.
- The first interface in the group must be 1, 5, 9, 13, 17, 21, 25, or 29. You can specify fc5 through fc8 but not fc7 through fc10.
- The groups of four interfaces do not need to be consecutive. You can specify fc1 through fc8 and fc17 through fc20.

To configure a Fibre Channel flow using Fabric Manager, follow these steps:

Step 1 Expand **End Devices** and then select **SSM Features** in the Physical Attributes pane.

You see the Intelligent Storage Services configuration, showing the FCWA tab in the Information pane shown in Figure 49-5.

Figure 49-5 Intelligent Storage Services Configuration FCWA Tab

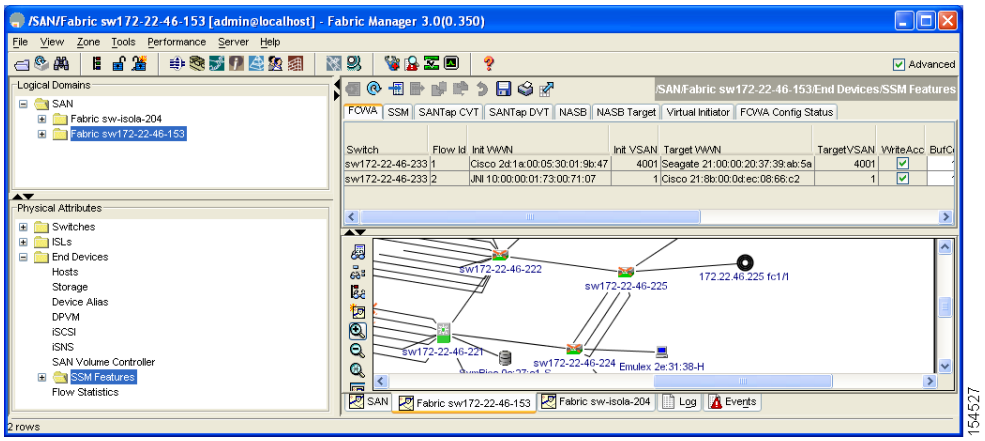

Step 2 Click **Create Row** in the Information pane to create a SCSI flow or click a row in the FCWA table to modify an existing SCSI flow.

You see the Fibre Channel write acceleration dialog box shown in Figure 49-6.

Figure 49-6 Fibre Channel Write Acceleration Dialog Box

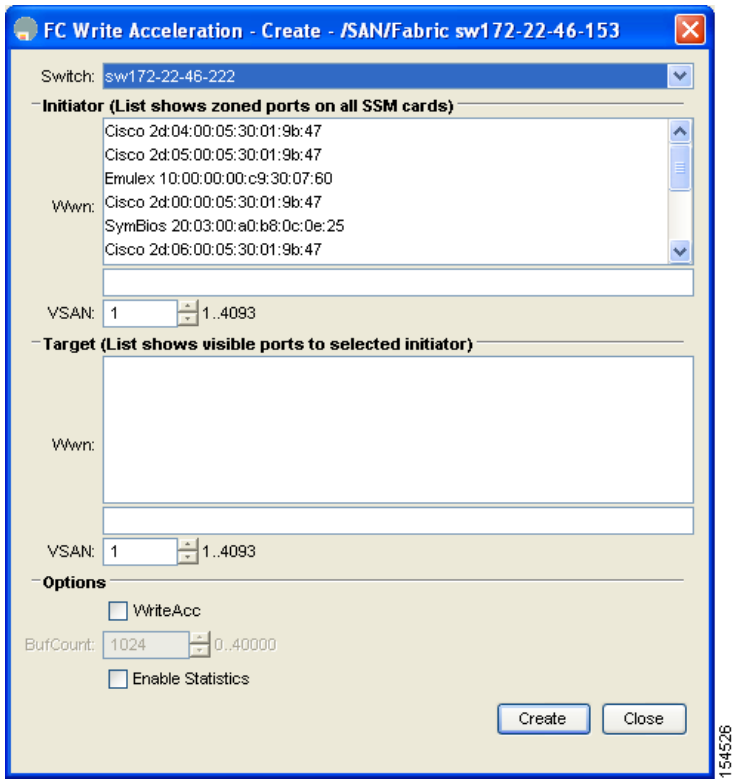

Step 3 Select the initiator and target WWNs and VSAN IDs and check the **WriteAcc** check box to enable Fibre Channel write acceleration on this SCSI flow.

Step 4 Optionally, enable SCSI flow statistics on this SCSI flow by checking the **Enable Statistics** check box.

Step 5 Optionally, change the BufCount value to set the number of 2K buffers used by the SCSI target.

Step 6 Click **Create** to create this SCSI flow or click **Cancel** to cancel this change.

About SCSI Flow Statistics

The statistics that can be collected for SCSI flows include the following:

- SCSI reads
  - Number of I/Os
  - Number of I/O blocks
  - Maximum I/O blocks
  - Minimum I/O response time
  - Maximum I/O response time
- SCSI writes
  - Number of I/Os
  - Number of I/O blocks
  - Maximum I/O blocks
  - Minimum I/O response time
  - Maximum I/O response time
- Other SCSI commands (not read or write)
  - Test unit ready
  - Report LUN
  - Inquiry
  - Read capacity
  - Mode sense
  - Request sense
- Errors
  - Number of timeouts
  - Number of I/O failures
  - Number of various SCSI status events
  - Number of various SCSI sense key errors or events

To take advantage of this feature, only the initiator must be directly attached to an SSM.

Note The SCSI flow statistics feature requires the Enterprise Package license installed only on the initiator switches.

Note For SCSI flow statistics, the initiator must connect to an SSM on a Cisco MDS switch while the target can connect to any other switch in the fabric. The SCSI flow initiator and target cannot connect to the same switch.

Enabling SCSI Flow Statistics

To enable SCSI flow statistics monitoring using Fabric Manager, follow these steps

Step 1 Expand **End Devices** and then select **SSM Features** in the Physical Attributes pane.

You see the Intelligent Storage Services configuration, showing the FCWA tab in the Information pane.

Step 2 Click **Create Row** in the Information pane to create a SCSI flow or click a row in the FCWA table to modify an existing SCSI flow.

You see the Fibre Channel write acceleration dialog box shown in Figure 49-7.

Figure 49-7 Fibre Channel Write Acceleration Dialog Box

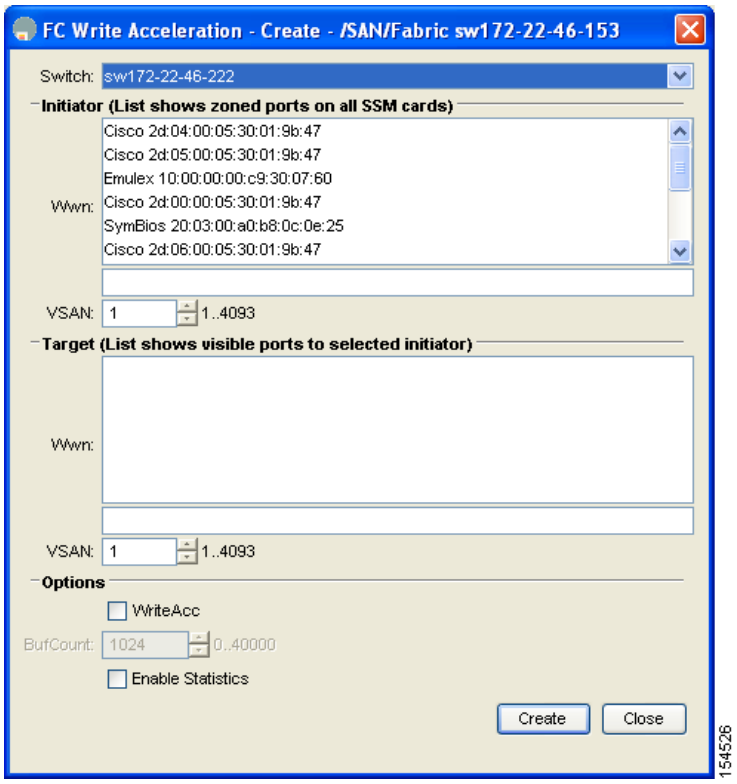

Step 3 Select the initiator and target WWNs and VSAN IDs and check the **Enable Statistics** check box to enable SCSI flow statistics on this SCSI flow.

Step 4 Optionally, enable Fibre Channel write acceleration on this SCSI flow at this time by checking the **WriteAcc** check box.

Step 5 Click **Create** to create this SCSI flow or click **Cancel** to cancel this change.

Clearing SCSI Flow Statistics

To clear SCSI flow statistics using Fabric Manager, follow these steps:

Step 1 Expand **End Devices** and then select **SSM Features**.

You see the FCWA tab information.

Step 2 Scroll the Information pane so you see the right half of the screen as shown in Figure 49-8.

Figure 49-8 Intelligent Storage Services Configuration FCWA Tab - Right Half

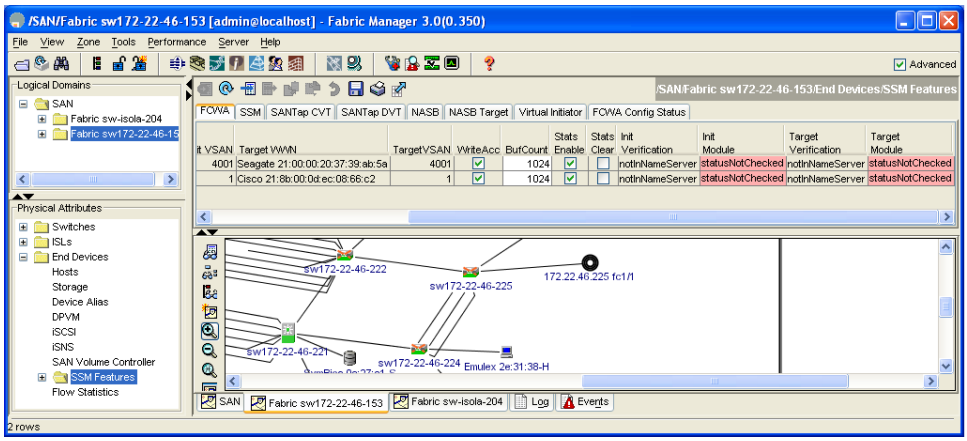

Step 3 Check the **Stats Clear** check box to clear SCSI flow statistics.

Step 4 Click **Apply Changes** to clear the SCSI flow statistics or click **Undo Changes** to discard any unsaved changes.

Default Settings

Table 49-1 lists the default settings for Intelligent Storage Services parameters.

Table 49-1 Default Intelligent Storage Services Parameters

Parameters	Default
SCSI flow services	Disabled.
SCSI flow services distribution	Enabled.
SCSI flow statistics	Disabled.

CHAPTER 50

Configuring Fibre Channel Write Acceleration

Intelligent Storage Services are features supported on the Storage Services Module (SSM). Intelligent Storage Services supported in Cisco MDS SAN-OS Release 2.0(2b) and later, including Fibre Channel write acceleration.

This chapter includes the following sections:

Intelligent Storage Services

This section contains the following topics:

Enabling Intelligent Storage Services

You can provision a subset of the ports for an SSM feature. The port range must be a multiple of four (for example fc4/1 through fc4-12).

To enable Intelligent Storage Services for an SSM and provision all ports or a group of ports to use these services using Fabric Manager, follow these steps:

Step 1 Expand **End Devices** and then select **SSM Features** in the Physical Attributes pane.

You see the Intelligent Storage Services configuration in the Information pane.

Step 2 Click the **SSM** tab.

You see the set of configured services in the Information pane shown in Figure 50-1.

Figure 50-1 SSM Configured Services

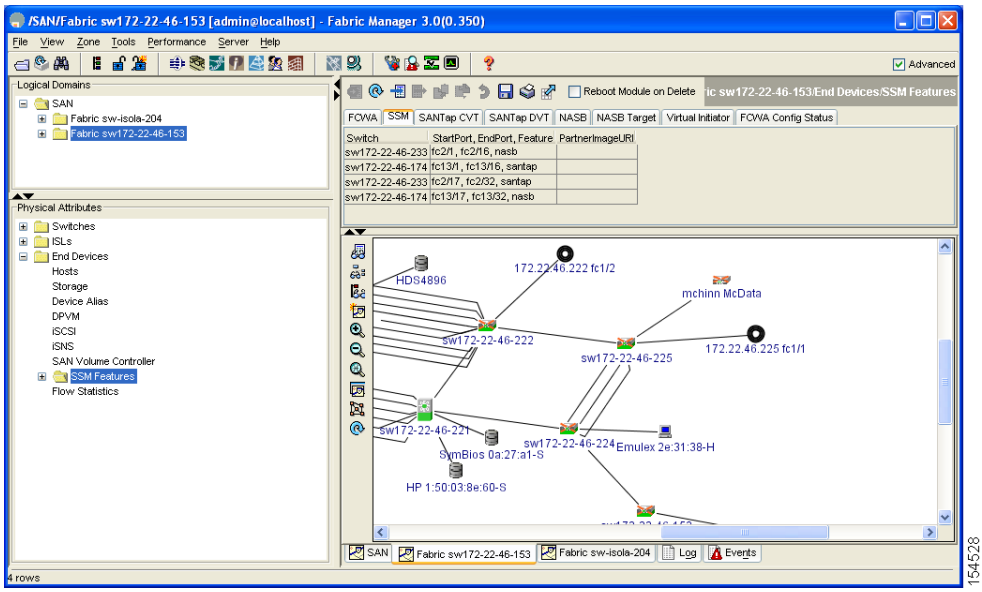

Step 3 Click **Create Row** to enable a new service on an SSM.

You see the Create SSM dialog box shown in Figure 50-2.

Figure 50-2 Create SSM Dialog Box

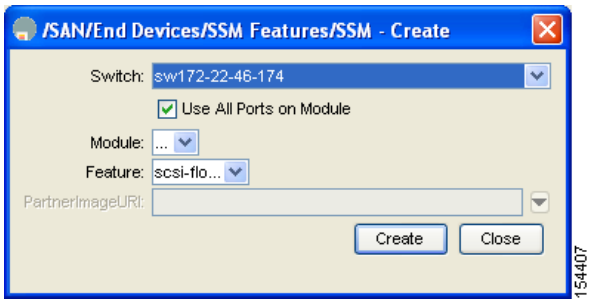

Step 4 Select the switch and SSM card you want to configure.

Step 5 Optionally, uncheck the **Use All Ports on Module** check box if you want to provision a subset of the ports on the card to use this service.

Step 6 Select the port range you want to provision for using this service (starting port and ending port).

Note The port range must be a multiple of four (for example fc4/1 through fc4-12).

Step 7 Select the feature you want to enable on these ports from the drop-down list of services.

Step 8 Set the PartnerImageURI field if you are enabling a third-party application that requires an image loaded onto the SSM.

Step 9 Click **Create** to create this row and enable this service or click **Cancel** to discard all changes.

Disabling Intelligent Storage Services

To disable Intelligent Storage Services in Fabric Manager for an SSM and free up a group of ports that used these services, follow these steps:

Step 1 Expand **End Devices** and then select **SSM Features** in the Physical Attributes pane.

You see the Intelligent Storage Services configuration in the Information pane.

Step 2 Click the **SSM** tab.

You see the set of configured services in the Information pane shown in Figure 50-3.

Figure 50-3 SSM Configured Services

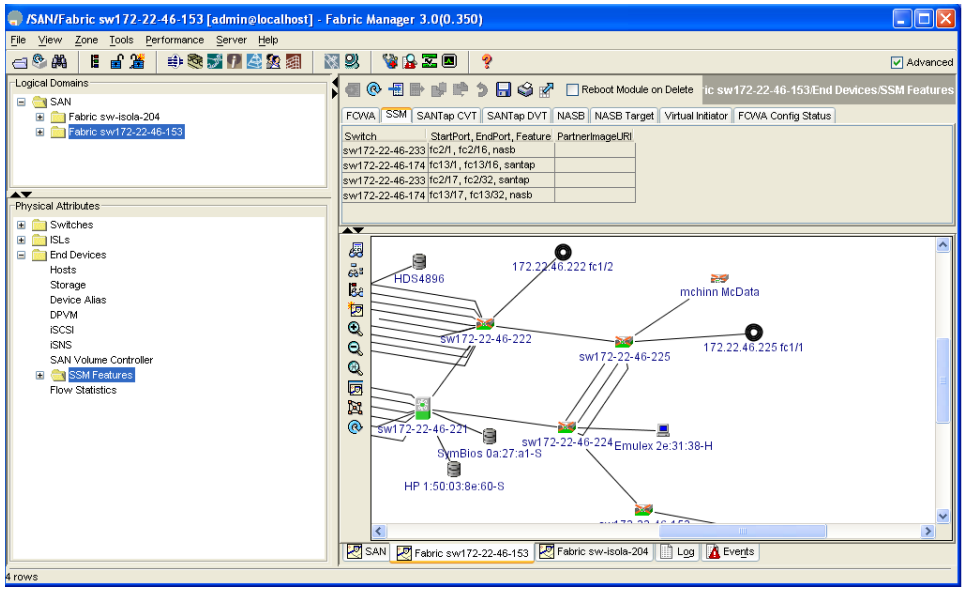

Step 3 Select the row in the table that you want to disable.

Step 4 Optionally, check the **Reboot Module on Delete** check box if you want to force the card to reboot after disabling the service. This is equivalent to the CLI **force** option.

Step 5 Click **Delete Row** to delete this row.

The ports that were provisioned for this service become available for provisioning in another service.

Note If **Reboot Module on Delete** was checked, then the SSM module reboots.

Fibre Channel Write Acceleration

Fibre Channel write acceleration minimizes application latency or reduces transactions per second over long distances. For synchronous data replication, Fibre Channel write acceleration increases the distance of replication or reduces effective latency to improve performance. To take advantage of this feature, both the initiator and target devices must be directly attached to an SSM.

This section includes the following topics:

About Fibre Channel Write Acceleration

The Fibre Channel write acceleration feature also allows the configuration of the buffer count. You can change the number of 2-KB buffers reserved on the target side DPP for a SCSI flow.

You can estimate the number of buffers to configure using the following formula:

(Number of concurrent SCSI writes * size of SCSI writes in bytes) / FCP data frame size in bytes

For example, for HDS TrueCopy between HDS 9970s, which use 1-KB FCP data frames, and you perform an initial sync for a 16-LUN TrueCopy group with 15 tracks, or 768-KB per LUN, the approximate number of write buffers required would be 16*(768*1024)/1024 or 12248 buffers.

Note The Fibre Channel write acceleration feature requires the Enterprise Package license installed on both the initiator and target switches.

Note The initiator and target cannot connect to the same Cisco MDS switch. Fibre Channel write acceleration requires that the initiator and target must connect to an SSM module on different Cisco MDS switches.

Enabling Fibre Channel Write Acceleration

To enable Fibre Channel write acceleration, and optionally modify the number of write acceleration buffers with Fabric Manager, follow these steps

Step 1 Expand **End Devices** and then select **SSM Features** from the Physical Attributes pane.

You see the Intelligent Storage Services configuration, showing the FCWA tab in the Information pane.

Figure 50-4 Intelligent Storage Services Configuration FCWA Tab

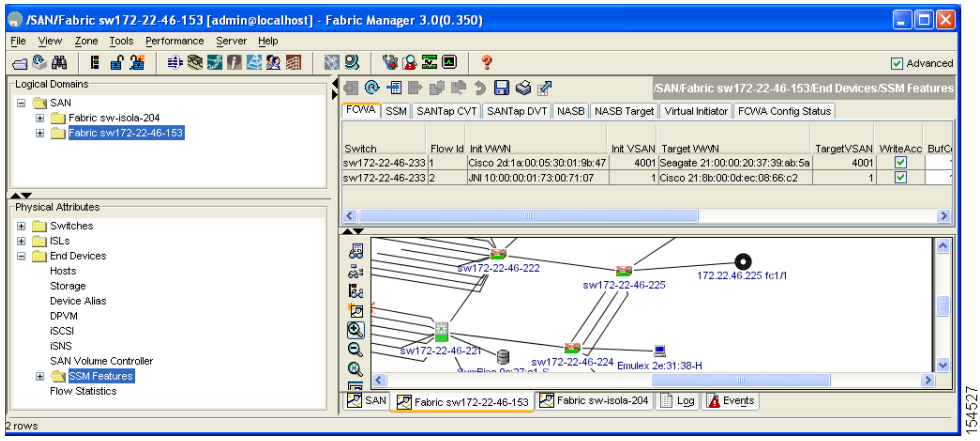

Step 2 Click **Create Row** in the Information pane to create a SCSI flow or click a row in the FCWA table to modify an existing SCSI flow.

You see the Fibre Channel write acceleration dialog box shown in Figure 50-5.

Figure 50-5 Fibre Channel Write Acceleration Dialog Box

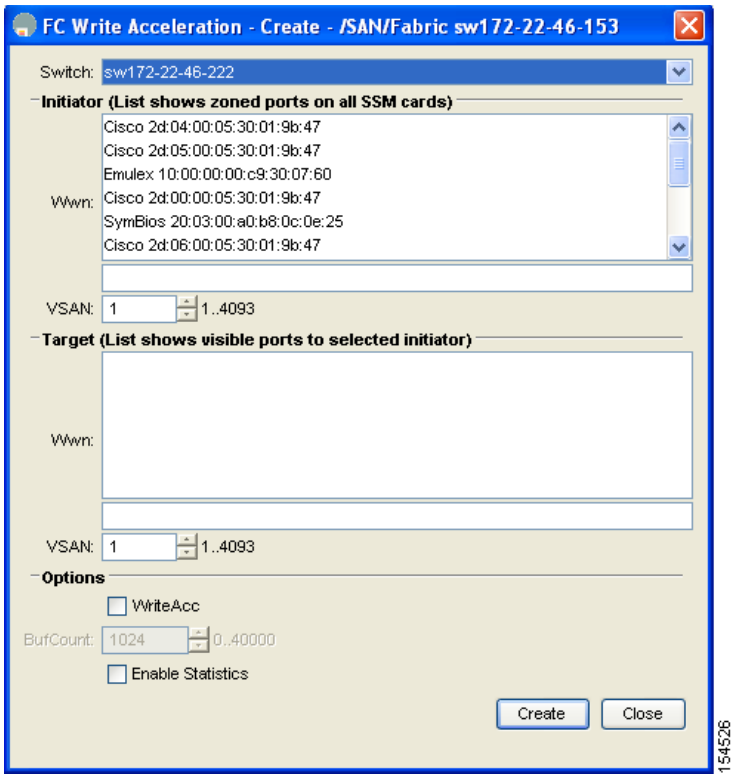

Step 3 Select the initiator and target WWNs and VSAN IDs and check the **WriteAcc** check box to enable Fibre Channel write acceleration on this SCSI flow.

Step 4 Optionally, enable SCSI flow statistics on this SCSI flow at this time by checking the **Enable Statistics** check box.

Step 5 Optionally, set the BufCount value to the number of 2K buffers used by the SCSI target.

Step 6 Click **Create** to create this SCSI flow with Fibre Channel write acceleration or click **Cancel** to discard all changes.

Default Settings

Table 50-1 lists the default settings for Fibre Channel write acceleration parameters.

Table 50-1 Default Intelligent Storage Services Parameters

Parameters	Default
Fibre Channel write acceleration	Disabled.
Fibre Channel write acceleration buffers	1024.

CHAPTER 51

Configuring SANTap

SANTap is one of the Intelligent Storage Services features supported on the Storage Services Module (SSM). The Storage Services Module (SSM) supports SANTap in Cisco MDS SAN-OS Release 2.0(2b) and later.

For licensing details, see Chapter 10, "Obtaining and Installing Licenses."

This chapter includes the following sections:

Intelligent Storage Services

All Intelligent Storage Services must be enabled on an SSM before the service can be configured. For switches running Cisco MDS SAN-OS Release 2.1(1a) or later software, some services are enabled for all ports on the SSM, or provisioned in groups of four ports. Switches running earlier releases that support intelligent storage services enable a service across all ports.

Note The four port groups are contiguous, requiring you to configure ports 1 through 4, 5 through 8, and so on.

This section includes the following topics:

Enabling Intelligent Storage Services

In Cisco MDS SAN-OS Release 2.1(1a) or later, you can provision a subset of the ports for an SSM feature. The port range must be a multiple of four (for example fc4/1 through fc4-12).

To enable Intelligent Storage Services for an SSM and provision all ports or a group of ports to use these services using Fabric Manager, follow these steps:

Step 1 Expand **Switches** and **End Devices** and then select **SSM Features** in the Physical Attributes pane.

You see the Intelligent Storage Services configuration in the Information pane.

Figure 51-1 SSM Features

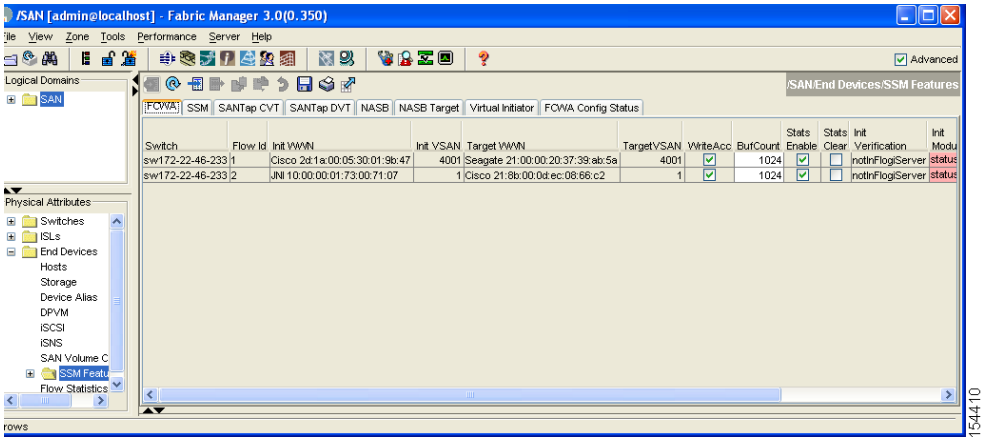

Step 2 Click the **SSM** tab.

You see the set of configured services in the Information pane.

Step 3 Click the **Create Row** icon to enable a new service on an SSM.

You see the SSM Create dialog box shown in Figure 51-2.

Figure 51-2 Create SSM Dialog Box

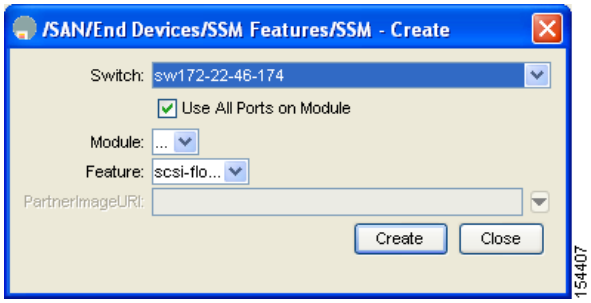

Step 4 Select the switch and SSM card you want to configure.

Step 5 Uncheck the **Use All Ports on Card** check box if you want to provision a subset of the ports on the card to use this service.

Step 6 Select the port range you want to provision for using this service (starting port and ending port.

Note The port range must be a multiple of four (for example fc4/1 through fc4-12).

Step 7 Select the Feature you want to enable on these ports from the drop-down list of services. The services that you can enable on a switch are SCSI Flow, EMCSR, NSP, SANTap, and NASB.

Step 8 Set the PartnerImageURI field if you are enabling a third-party application that requires an image loaded onto the SSM.

Step 9 Click **Create** to create this row and enable this service or click **Cancel** to discard all changes.

Disabling Intelligent Storage Services

To disable Intelligent Storage Services in Fabric Manager for an SSM and free up a group of ports that used these services, follow these steps:

Step 1 Expand **Switches** and **End Devices** and then select **SSM Features** in the Physical Attributes pane.

You see the Intelligent Storage Services configuration in the Information pane.

Figure 51-3 SSM Features

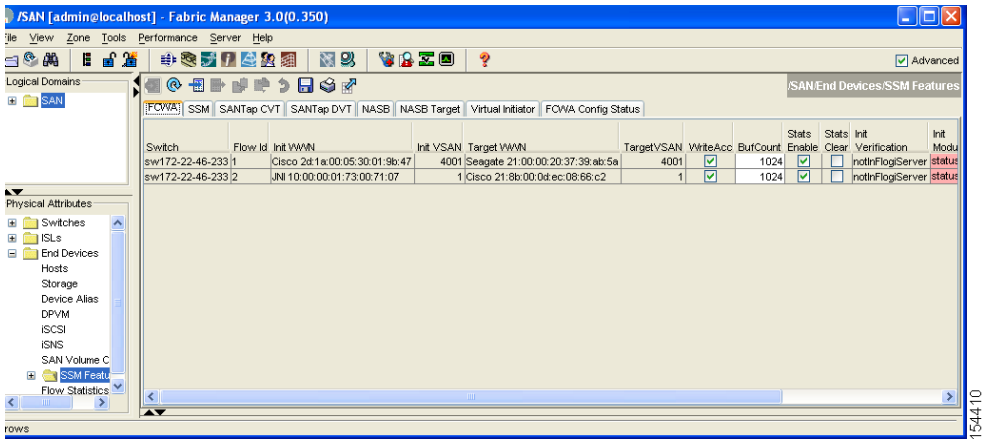

Step 2 Click the **SSM** tab.

You see the set of configured services in the Information pane.

Step 3 Select the row in the Configured Services table that you want to disable.

Step 4 Check the **Reboot Module on Delete** check box if you want to force the card to reboot after disabling the service. This is equivalent to the CLI force option.

Step 5 Click the Delete Row icon to delete this row. The ports that were provisioned for this service become available for provisioning in another service.

Note If the **Reboot Card on Delete** check box was checked, then the SSM module reboots.

About SANTap

The SANTap feature allows third-party data storage applications, such as long distance replication and continuous backup, to be integrated into the SAN. The protocol-based interface that is offered by SANTap allows easy and rapid integration of the data storage service application because it delivers a loose coupling between the application and an SSM, thereby reducing the effort needed to integrate applications with the core services being offered by the SSM. See Figure 51-4.

Figure 51-4 Integrating Third-Party Storage Applications in a SAN

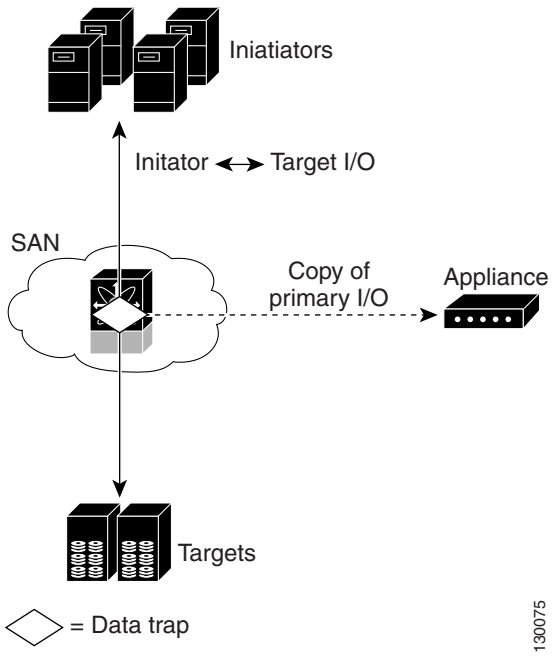

SANTap has a control path and a data path. The control path services requests that create and manipulate replication sessions sent by an appliance. The control path is implemented using a SCSI-based protocol. An appliance sends requests to a control virtual target (CVT) which the SANTap process creates and monitors. Responses are sent to the control LUN on the appliance. SANTap also allows LUN mapping to appliance virtual targets (AVTs). You can have a maximum of 512 target LUNs.

SANTap does not require reconfiguration of either the host or target when introducing SANTap-based applications. Also, neither the host initiator nor the target is required to be directly connected to an SSM. This accomplished by assigning Cisco-specific WWNs to the virtual initiators (VIs) and digital virtual targets (DVTs). A host initiator or a target be connected directly to an SSM. However, you must partition the SAN using VSANs.

You must configure the host initiator and the DVT in one VSAN and configure the virtual initiator (VI) and the target in another VSAN. See Figure 51-5.

Figure 51-5 SANTap Proxy Mode-2 Example

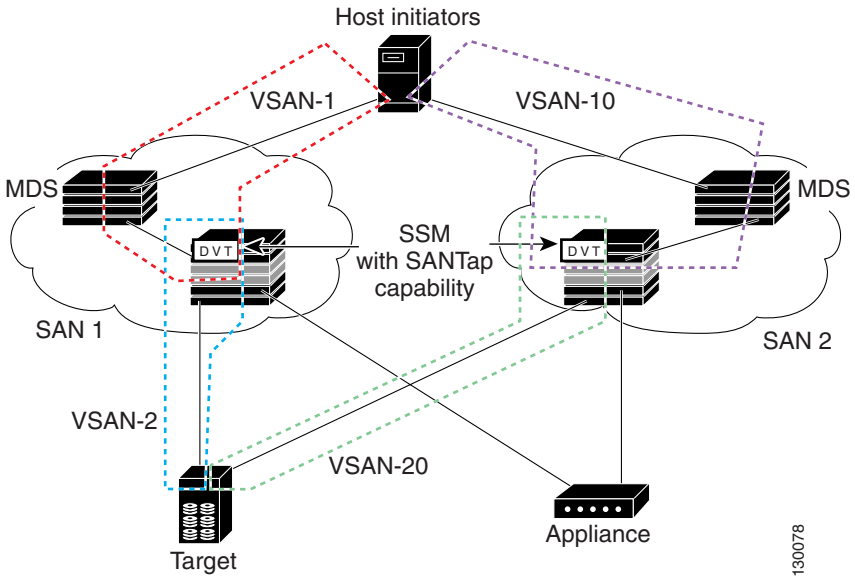

This section includes the following topics:

About Enabling SANTap

SANTap can be enabled on an entire SSM or it can be enabled on a group of four ports on an SSM. Enabling SANTap on interfaces has the following restrictions:

- The fewest number of interfaces that you can enable is four. You can specify fc1 through fc4 but not fc1 through fc2.
- The first interface in the group must be 1, 5, 9, 13, 17, 21, 25, or 29. You can specify fc5 through fc8 but not fc7 through fc10.
- The groups of four interfaces do not need to be consecutive. You can specify fc1 through fc8 and fc17 through fc20.

Enabling SANTap

See "Enabling Intelligent Storage Services" section on page 51-1. In the Create SSM dialog, select SANTap from the Feature drop-down list.

Creating a SANTap CVT

It is required to configure a logical port on a switch to create the Control Virtual Target (CVT) for SANTap. CVTs create the control path which processes the SANTap service requests sent by an appliance.

Before requesting the SANTap service the appliance contacts the CVT, specifies the initiator and the target for replicating the data flowing between them.

To create a SANTap Control Virtual Target (CVT) for SANTap on the Fibre Channel switch modules using Fabric Manager, follow these steps:

Step 1 Expand **Switches > End Devices** and then select **SSM Features** from the Physical Attributes pane.

You see the Intelligent Storage Services configuration showing the FCWA tab, in the Information pane.

Step 2 Click the **SANTap CVT** tab.

You see the SANTap configuration in the Information pane.

Step 3 Click **Create Row**.

You see the create SANTap CVT configuration dialog box shown in Figure 51-6.

Figure 51-6 Create SANTap CVT Configuration Dialog Box

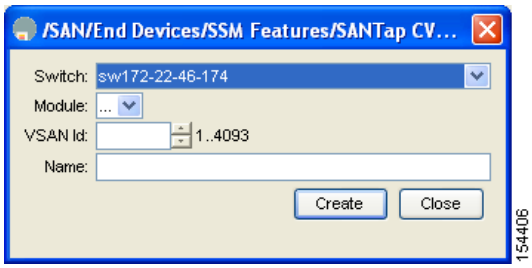

Step 4 Select the Switch and the Module on which you want to configure a SANTap CVT.

Note SANTap must be enabled and provisioned as a service on the SSM module of the selected switch. See the "Enabling Intelligent Storage Services" section on page 51-1.

Step 5 Select the VSAN ID in which you want to configure the SANTap CVT.

Step 6 Click **Create** to create this SANTap CVT or click **Cancel**.

Deleting a SANTap CVT

To delete a SANTap CVT using Fabric Manager, follow the steps below:

Step 1 Expand **Switches > End Devices** and then select **SSM Features** from the Physical Attributes pane.

You see the Intelligent Storage Services configuration showing the FCWA tab, in the Information pane.

Step 2 Click the **SANTap CVT** tab.

You see the SANTap configuration in the Information pane.

Step 3 Select the SANTap CVT you want to delete.

Step 4 Click **Delete Row**.

You see the Fabric Manager Confirmation dialog box.

Step 5 Click **Yes**, to proceed with the deletion or click **No** to discard the changes.

Creating a SANTap DVT

To create a SANTap Destination Virtual Target (DVT) using Fabric Manager, follow the steps below:

Step 1 Expand **Switches > End Devices** and then select **SSM Features** from the Physical Attributes pane.

You see the Intelligent Storage Services configuration showing the FCWA tab, in the Information pane.

Step 2 Click the **SANTap DVT** tab.

You see the SANTap configuration in the Information pane.

Figure 51-7 SANTap DVT Tab

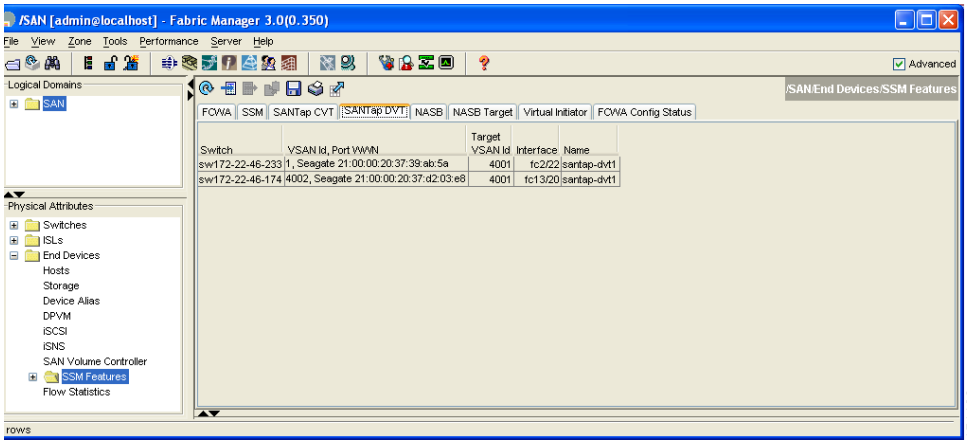

Step 3 Click **Create Row**.

You see the create SANTap DVT configuration dialog box shown in Figure 51-8.

Figure 51-8 Create SANTap DVT Configuration Dialog Box

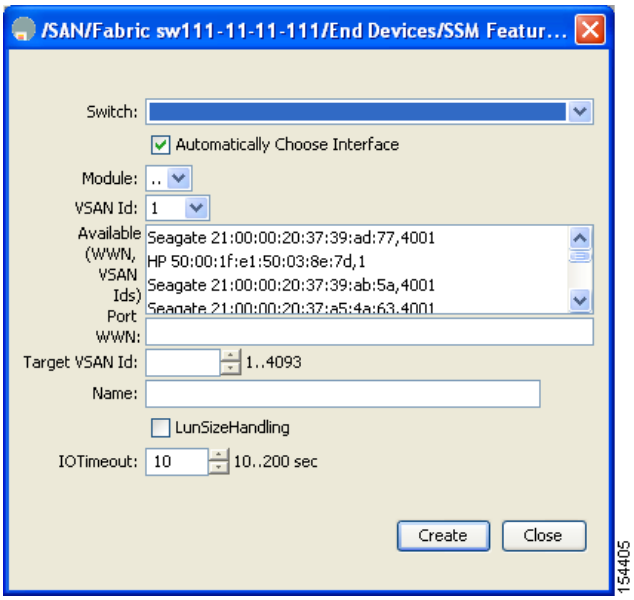

Step 4 Select the switch on which the SANTap DVT will be configured.

Step 5 Select the VSAN ID in which you want to create the SANTap DVT.

Step 6 Select the port WWN of the real target for which this corresponding DVT is being created. The DVT has the same port WWN as the target.

Step 7 Select the target VSAN ID for the VSAN of the real target for which this DVT is being created.

Step 8 Select the interface. This is the port on the module where the DVT will be created.

Step 9 Assign a name to this SANTap DVT.

Step 10 Check the LunSizeHandling option if you want to use the real target LUN size for the virtual LUN or the maximum LUN size supported (2TB).

Step 11 Select the IOTimeout value for the DVT. The default value is 10 seconds.

Step 12 Click **Create** to create this SANTap DVT or click **Cancel** to discard the changes.

Deleting a SANTap DVT

To delete a SANTap DVT using Fabric Manager, follow the steps below:

Step 1 Expand **Switches > End Devices** and then select **SSM Features** from the Physical Attributes pane.

You see the Intelligent Storage Services configuration showing the FCWA tab, in the Information pane.

Step 2 Click the **SANTap DVT** tab.

You see the SANTap configuration in the Information pane.

Step 3 Select the SANTap DVT you want to delete.

Step 4 Click **Delete Row**.

You see the Fabric Manager Confirmation dialog box.

Step 5 Click **Yes** to proceed with the deletion or click **No** to discard the changes.

Default Settings

Table 51-1 lists the default settings for Intelligent Storage Services parameters.

Table 51-1 Default SANTap Parameters

Parameters	Default
SANTap feature	Disabled.

CHAPTER 52

Configuring NASB

Intelligent Storage Services are features supported on the Storage Services Module (SSM). Intelligent Storage Services supported in Cisco MDS SAN-OS Release 2.1(1a) and later include Network-Accelerated Serverless Backup (NASB).

For licensing details, see Chapter 10, "Obtaining and Installing Licenses."

This chapter includes the following sections:

Intelligent Storage Services

All Intelligent Storage Services must be enabled on an SSM before the service can be configured. For switches running Cisco MDS SAN-OS Release 2.1(1a) or later software, some services are enabled for all ports on the SSM, or provisioned in groups of four ports. Switches running earlier releases that support intelligent storage services enable a service across all ports.

Note The four port groups are contiguous, requiring you to configure ports 1 through 4, 5 through 8, and so on.

This section contains the following topics:

Enabling Intelligent Storage Services

In Cisco MDS SAN-OS Release 2.1(1a) or later, you can provision a subset of the ports for an SSM feature. The port range must be a multiple of four (for example fc4/1 through fc4-12).

To enable Intelligent Storage Services for an SSM and provision all ports or a group of ports to use these services using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **End Devices** and then select **SSM Features** in the Physical Attributes pane.

You see the Intelligent Storage Services configuration in the Information pane. See Figure 52-1.

Figure 52-1 SSM Features

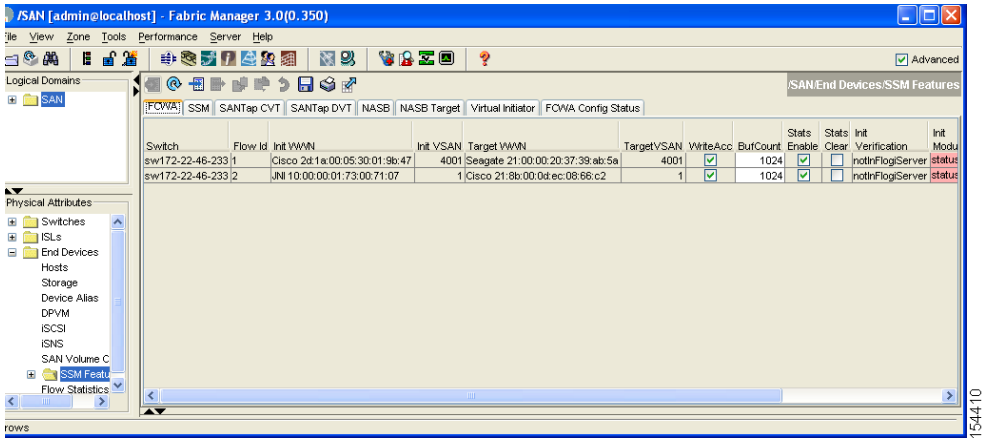

Step 2 Click the **SSM** tab.

You see the set of configured services in the Information pane. See Figure 52-2.

Figure 52-2 SSM Configured Services

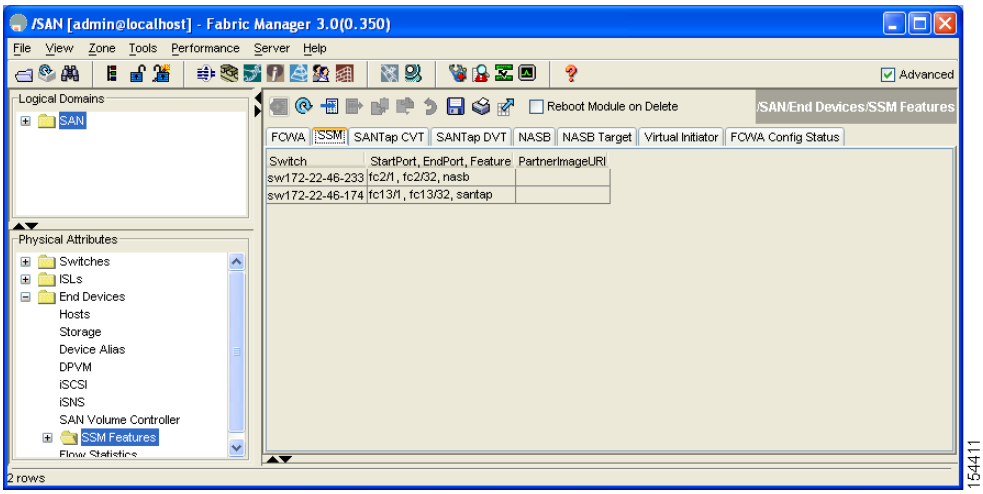

Step 3 Click **Create Row** to enable a new service on an SSM.

You see the Create SSM dialog box shown in Figure 52-3.

Figure 52-3 Create SSM Dialog Box

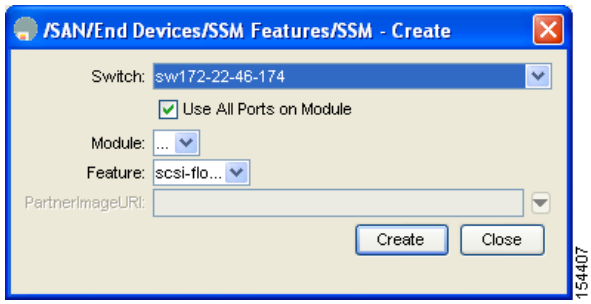

Step 4 Select the Switch and SSM Card you want to configure.

Step 5 Uncheck the **Use All Ports on Card** check box if you want to provision a subset of the ports on the card to use this service.

Step 6 Select the port range you want to provision for using this service (starting port and ending port).

Note The port range must be a multiple of four (for example fc4/1 through fc4-12).

Step 7 Select the Feature you want to enable on these ports from the drop-down list of services.

Step 8 Set the PartnerImageURI field if you are enabling a third-party application that requires an image loaded onto the SSM.

Step 9 Click **Create** to create this row and enable service or click **Cancel** to discard all changes.

Disabling Intelligent Storage Services

To disable Intelligent Storage Services in Fabric Manager for an SSM and free up a group of ports that used these services, follow these steps:

Step 1 Expand **Switches**, expand **End Devices** and then select **SSM Features** in the Physical Attributes pane.

You see the Intelligent Storage Services configuration in the Information pane.

Step 2 Click the **SSM** tab.

You see the set of configured services in the Information pane.

Step 3 Select the configured service in the table that you want to disable.

Step 4 Check the **Reboot Module on Delete** check box if you want to force the card to reboot after disabling the service. This is equivalent to the CLI force option.

Step 5 Click **Delete Row**. The ports that were provisioned for this service become available for provisioning in another service.

Note If the **Reboot Card on Delete** checkbox was checked, then the SSM module reboots.

NASB

Data movement in the fabric uses considerable processor cycles, which can cause client applications to slow down noticeably. Offloading data movement operations to a media server allows the client applications to run normally even during a backup operation. Media servers can further offload the data movement operation to NASB devices, which allows the media server to focus on the coordination functions needed to complete the backup.

Most backups performed today are server-free. In server-free backups, the application server is not involved in moving the data. The data can be moved by either a media server or a NASB device.

When the media server is the data mover, it moves the data between the disks and the tapes. The backup application runs on both the client device and the media server. However, the backup application in the client device performs minimal tasks for the backup operation.

The media server performs the following backup operations:

- Manages disks as well as one or more tape backup devices.
- Contacts the client devices to retrieve the list of logical blocks that need to be backed up.
- Performs data movement from disk to tape media based on the logical block list provided by the client device.

The backup application in the client device maps the data to be backed up and creates the logical block list associated with the data. The movement of data from the physical disks to the backup device (tape) is not performed by the client device. This reduces substantial load on the client device.

An example configuration is shown in Figure 52-4. The media server moves the data directly between the storage disks and the tape devices during backups.

Figure 52-4 Example Configuration with Media Server as Data Mover

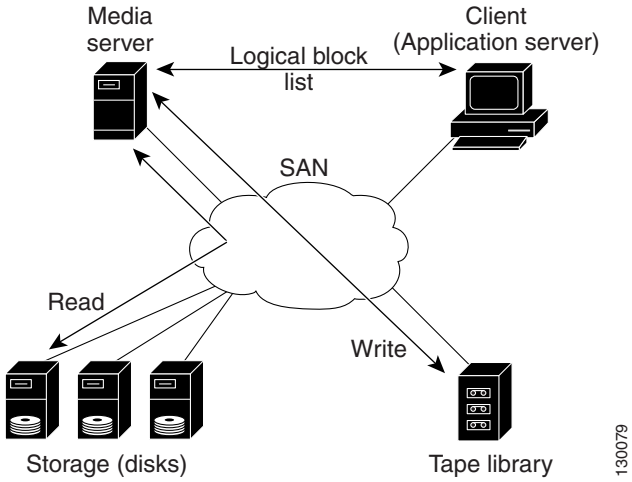

When the NASB is the data mover, it moves the data between the disks and the tapes. The NASB device is a SCSI target device capable of handling SCSI Extended Copy (XCOPY) commands as well as a SCSI initiator device capable of issuing READ/WRITE commands to disks and other backup media, such as tapes. See Figure 52-5.

Figure 52-5 Example Configuration with NASB Device as Data Mover

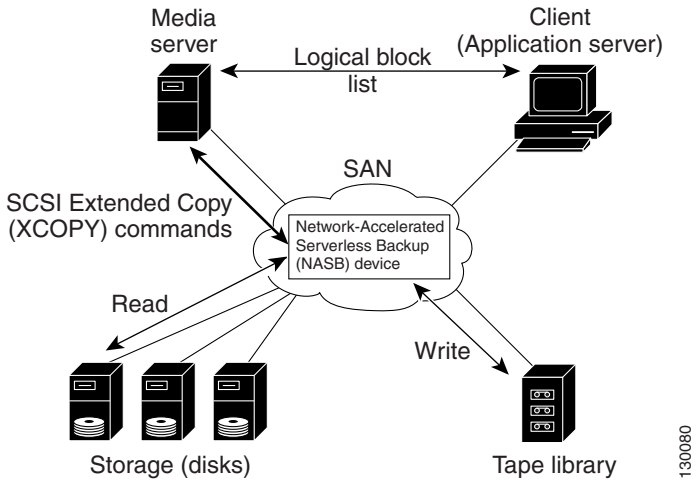

The task of managing and preparing the source and destination targets is performed by the media server. For example, if the destination is a tape library, the media server issues commands to load and unload the correct tape and position of the tape write head at the correct offset within the tape.

This section contains the following topics:

About Configuring NASB

Network-Accelerated Serverless Backup (NASB) can be enabled on an entire SSM or it can be enabled on one or more groups of four ports on an SSM. Enabling NASB on interfaces has the following restrictions:

- The fewest number of interfaces that you can enable is four. You can specify fc1 through fc4 but not fc1 through fc2.
- The first interface in the group must be 1, 5, 9, 13, 17, 21, 25, or 29. You can specify fc5 through fc8 but not fc7 through fc10.
- The groups of four interfaces do not need to be consecutive. You can specify fc1 through fc8 and fc17 through fc20.

Configuring NASB

To configure NASB using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **End Devices** and then select **SSM Features** in the Physical Attributes pane.

You see the Intelligent Storage Services configuration in the Information pane.

Step 2 Click the **NASB** tab.

You see the NASB configuration in the Information pane.

Step 3 Click **Create Row**.

You see the NASB configuration dialog box in Figure 52-6.

Figure 52-6 NASB Configuration Dialog Box

Step 4 Select the Switch and Module on which you want to configure NASB.

Note NASB must be enabled and provisioned as a service on this SSM module. See the "Enabling Intelligent Storage Services" section on page 52-2.

Step 5 Select the VSAN ID you want to configure for NASB.

Note You must configure this VSAN to permit default zoning.

Step 6 Choose the **Control** option to enable NASB on the SSM in one slot and on one VSAN for a single target LUN that is a Storage Array Controller (Peripheral Device Type = 0x0C).

Step 7 Choose the **Multiple** option to enable NASB on the SSM in one slot and on one VSAN for up to 10 target LUNs that are Direct Access Devices (Peripheral Device Type = 0x00).

Note Use the multiple option for multi-streaming (multiple backup sessions) on a single virtual target for VERITAS NetBackup.

Note For information on specific settings, consult with your NASB partner.

Step 8 Click **Create** to create this NASB or click **Cancel** to cancel this change.

Default Settings

Table 52-1 lists the default settings for Intelligent Storage Services parameters.

Table 52-1 Default Intelligent Storage Services Parameters

Parameters	Default
NASB feature	Disabled.

Cisco Systems

Send documentation comments to mdsfeedback-doc@cisco.com

PART 8

Network and Switch Monitoring

CHAPTER 53

Network Monitoring

The primary purpose of Fabric Manager is to manage the network. In particular, SAN discovery and network monitoring are two of its key network management capabilities.

This chapter contains the following sections:

SAN Discovery and Topology Mapping

Fabric Manager provides extensive SAN discovery, topology mapping, and information viewing capabilities. Fabric Manager collects information on the fabric topology through SNMP queries to the switches connected to it. Fabric Manager recreates a fabric topology, presents it in a customizable map, and provides inventory and configuration information in multiple viewing options.

Device Discovery

Once Fabric Manager is invoked, a SAN discovery process begins. Using information polled from a seed Cisco MDS 9000 Family switch, including Name Server registrations, Fibre Channel Generic Services (FC-GS), Fabric Shortest Path First (FSPF), and SCSI-3, Fabric Manager automatically discovers all devices and interconnects on one or more fabrics. All available switches, host bus adapters (HBAs), and storage devices are discovered. The Cisco MDS 9000 Family switches use Fabric-Device Management Interface (FMDI) to retrieve HBA model, serial number and firmware version, and host operating-system type and version discovery without host agents. Fabric Manager gathers this information through SNMP queries to each switch. The device information discovered includes device names, software revision levels, vendor, ISLs, PortChannels, and VSANs.

Topology Mapping

Fabric Manager is built upon a topology representation of the fabric. Fabric Manager provides an accurate view of multiple fabrics in a single window by displaying topology maps based on device discovery information. You can modify the topology map icon layout with an easy-to-use, drag-and-drop interface. The topology map visualizes device interconnections, highlights configuration information such as zones, VSANs, and ISLs exceeding utilization thresholds. The topology map also provides a visual context for launching command-line interface (CLI) sessions, configuring PortChannels, and opening device managers.

Using the Topology Map

The Fabric Manager topology map can be customized to provide a view into the fabric that varies from showing all switches, end devices, and links, to showing only the core switches with single bold lines for any multiple links between switches. Use the icons along the left side of the topology map to control these views or right-click anywhere in the topology map to access the map controls.

You can zoom in or out on the topology map to see an overview of the SAN or focus on an area of importance. You can also open an overview window that shows the entire fabric. From this window, you can right-click and draw a box around the area you want to view in the main topology map view.

Another way to limit the scope of the topology display is to select a fabric or VSAN from the Logical Domains pane. The topology map displays only that fabric or VSAN.

Moving the mouse pointer over a link or switch provides a simple summary of that SAN component, along with a status indication. Right-clicking on the component brings up a pop-up menu. You can view the component in detail or access configuration or test features for that component.

Double-click a link to bring link status and configuration information to the information pane. Double-click a switch to bring up Device Manager for that switch.

Saving a Customized Topology Map Layout

Changes made to the topology map can be saved so that the customized view is available any time you open the Fabric Manager Client for that fabric.

To save the customized layout using Fabric Manager, follow these steps:

Step 1 Click **File > Preferences** to open the Fabric Manager preferences dialog box.

Step 2 Click the **Map** tab and check the **Automatically Save Layout** check box to save any changes to the topology map. (See Figure 53-1.)

Figure 53-1 Fabric Manager Preferences

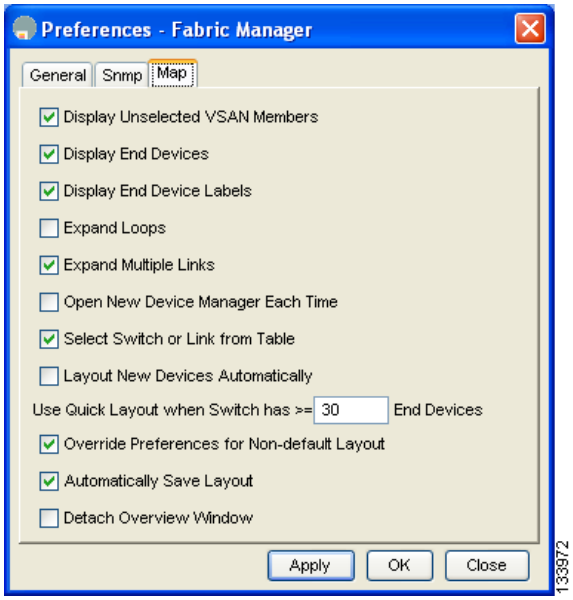

Step 3 Click **Apply** then **OK** to save this change, or click **Close** to discard any unsaved changes and close the dialog box.

Using Enclosures with Fabric Manager Topology Maps

Because not all devices are capable of responding to FC-GS-3 requests, different ports of a single server or storage subsystem may be displayed as individual end devices on the topology map. See the "Modifying the Device Grouping" section on page 4-17 to group these ports into a single enclosure for Fabric Manager.

Clicking **Alias->Enclosure** displays hosts and storage elements in the Information pane. This is a shortcut to naming enclosures. To use this shortcut, highlight each row in the host or storage table that you want grouped in an enclosure then click **Alias -> Enclosure**. This automatically sets the enclosure names of each selected row with the first token of the alias.

Mapping Multiple Fabrics

To log into multiple fabrics, the same username and password must be used. The information for both fabrics is displayed, with no need to select a seed switch. To see details of a fabric, click the tab for that fabric at the bottom of the Fabric pane, or double-click the fabric's cloud icon.

When you quit the Fabric Manager Client, you can have Fabric Manager Server continuously monitor that fabric. Alternatively, you can use Fabric Manager client to select a fabric to monitor.

To continuously monitor a fabric using Fabric Manager, follow these steps:

Step 1 Choose **Server > Admin**.

You see the Server Admin dialog box with a list of fabrics shown in Figure 53-2.

Figure 53-2 Server Admin Dialog Box

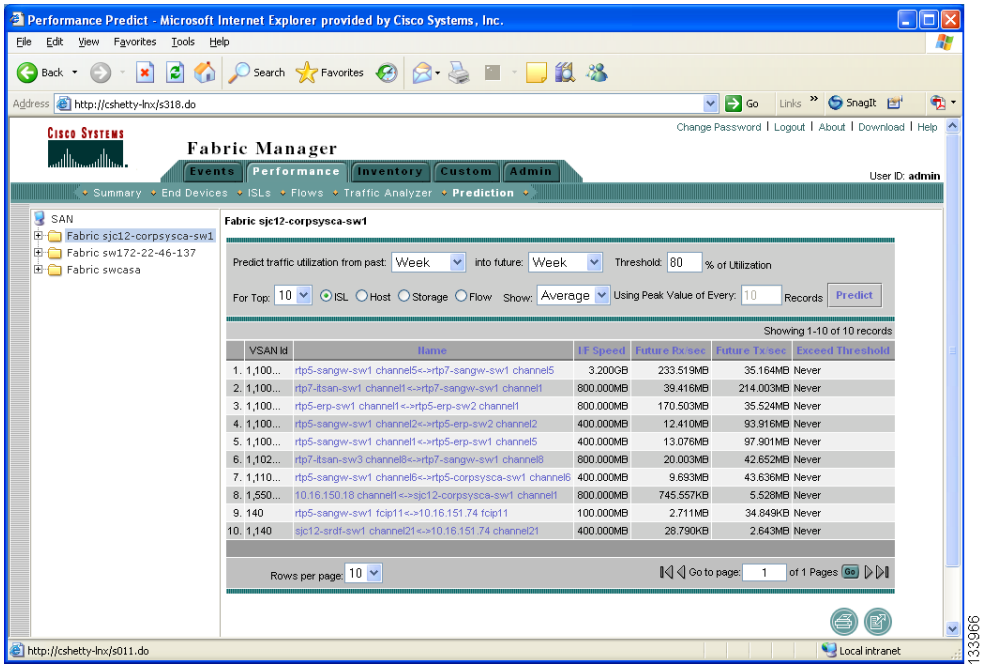

Step 2 Check the **Monitor Continuously** check box of the fabric(s) you want Fabric Manager Server to monitor.

Step 3 Click **Apply.**

The Monitor Continuously feature requires the purchase of the Fabric Manager Server license package. If you have not purchased and installed this package, you see a pop-up window informing you that you are about to enable a demo license for this feature. Click **Yes** to enable the demo license.

Note When you are finished checking out the demo, you can "check in" the feature by clicking **Check In FM** as described in the "Fabric Manager Server Licensing" section on page 10-16.

Step 4 Click **Close** to close the Server Admin dialog box.

Inventory Management

The Information pane in Fabric Manager shows inventory, configuration, and status information for all switches, links, and hosts in the fabric. Inventory management includes vendor name and model, and software or firmware versions. Select a fabric or VSAN from the Logical Domains pane, and then select the **Summary** tab in the Information pane to get a count of the number of VSANS, switches, hosts, and storage elements in the fabric. See the "Fabric Manager Client Quick Tour" section on page 4-3 for more information on the Fabric Manager user interface.

Using the Inventory Tab from Fabric Manager Web Services

If you have configured Fabric Manager Web Services, you can launch this application and access the Inventory tab to see a summary of the fabrics managed by the Fabric Manager Server. The Inventory tab shows an inventory of the selected SAN, fabric, or switch.

- Summary—Shows all VSANs, switches, and ports in the selected SAN or fabric.
- VSANs— Shows all VSANs in the selected SAN or fabric.
- Switches—Shows all attributes (such as IP address, vendor, and model) for all switches in the selected SAN, fabric, or VSAN.
- Licenses—Shows details about the licenses in use in the fabric.
- Modules—Shows all line cards, fans, and power supplies for all switches in the selected SAN, fabric, or VSAN.
- End Devices—Shows the host and storage ports.
- ISLs—Shows all the Inter-Switch Links for the selected SAN, fabric, or VSAN.
- Zones—Shows all the active zone members (including those in inter-VSAN zones) for the selected SAN, fabric, or VSAN.

See Chapter 6, “Fabric Manager Web Services” for more information on how to configure and use Fabric Manager Web Services.

To view system messages remotely using Fabric Manager Web Services, follow these steps:

Step 1 Point your browser at the Fabric Manager Web Services server. See the “Launching Fabric Manager Web Services” section on page 6-5.

Step 2 Click the **Events** tab then the **Details** to view the system messages. The columns in the events table are sortable. In addition, you can use the Filter button to limit the scope of messages within the table.

Viewing Logs from Device Manager

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

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

Health and Event Monitoring

Fabric Manager works with the Cisco MDS 9000 Family switches to show the health and status of the fabric and switches. Information about the fabric and its components is gathered from multiple sources, including Online System Health Management, Call Home, system messages, and SNMP notifications. This information is then made available from multiple menus on Fabric Manager or Device Manager.

Fabric Manager Events Tab

The Fabric Manager Events tab, available from the topology window, displays the events Fabric Manager received from sources within the fabric. These sources include SNMP events, RMON events, system messages, and system health messages. The Events tab shows a table of events, including the event name, the source and time of the event, a severity level, and a description of the event. The table is sortable by any of these column headings.

Event Information in Fabric Manager Web Services Reports

The Fabric Manager web services client displays collections of information gathered by the Performance Manager. This information includes events sent to the Fabric Manager Server from the fabric. To open these reports, choose **Performance Manager > Reports**. This opens the web client in a web browser and displays a summary of all fabrics monitored by the Fabric Manager Server. Choose a fabric and then click the **Events** tab to see a summary or detailed report of the events that have occurred in the selected fabric. The summary view shows how many switches, ISLs, hosts, or storage elements are down on the fabric and how many warnings have been logged for that SAN entity. The detailed view shows a list of all events that have been logged from the fabric and can be filtered by severity, time period, or type.

Events in Device Manager

Device Manager displays the events when you choose **Logs > Events**. Device Manager can display the current list of events or an older list of events that has been stored on the Fabric Manager host. The event table shows details on each event, including time, source, severity, and a brief description of the event.

CHAPTER 54

Performance Monitoring

Cisco Fabric Manager and Device Manager provide multiple tools for monitoring the performance of the overall fabric, SAN elements, and SAN links. These tools provide real-time statistics as well as historical performance monitoring.

This chapter contains the following sections:

Real-Time Performance Monitoring

Real-time performance statistics are a useful tool in dynamic troubleshooting and fault isolation within the fabric. Real-time statistics gather data on parts of the fabric in user-configurable intervals and display these results in Fabric Manager and Device Manager.

Device Manager Real-Time Performance Monitoring

Device Manager provides an easy tool for monitoring ports on the Cisco MDS 9000 Family switches. This tool gathers statistics at a configurable interval and displays the results in tables or charts. These statistics show the performance of the selected port in real-time and can be used for performance monitoring and troubleshooting. For a selected port, you can monitor any of a number of statistics including traffic in and out, errors, class 2 traffic, and FICON data. You can set the polling interval from ten seconds to one hour, and display the results based on a number of selectable options including absolute value, value per second, and minimum or maximum value per second.

Device Manager in Cisco MDS SAN-OS Release 2.1(1) or later supports checking for oversubscription on the host-optimized four-port groups on relevant modules. Right-click the port group on a module and choose **Check Oversubscription** from the pop-up menu.

Device manager provides two performance views, the Summary View tab, and the configurable monitor option per port.

To configure the summary view in Device Manager, follow these steps:

Step 1 Click the **Summary** tab on the main display.

You see all of the active ports on the switch, as well as the configuration options available from the Summary view shown in Figure 54-1.

Figure 54-1 Device Manager Summary Tab

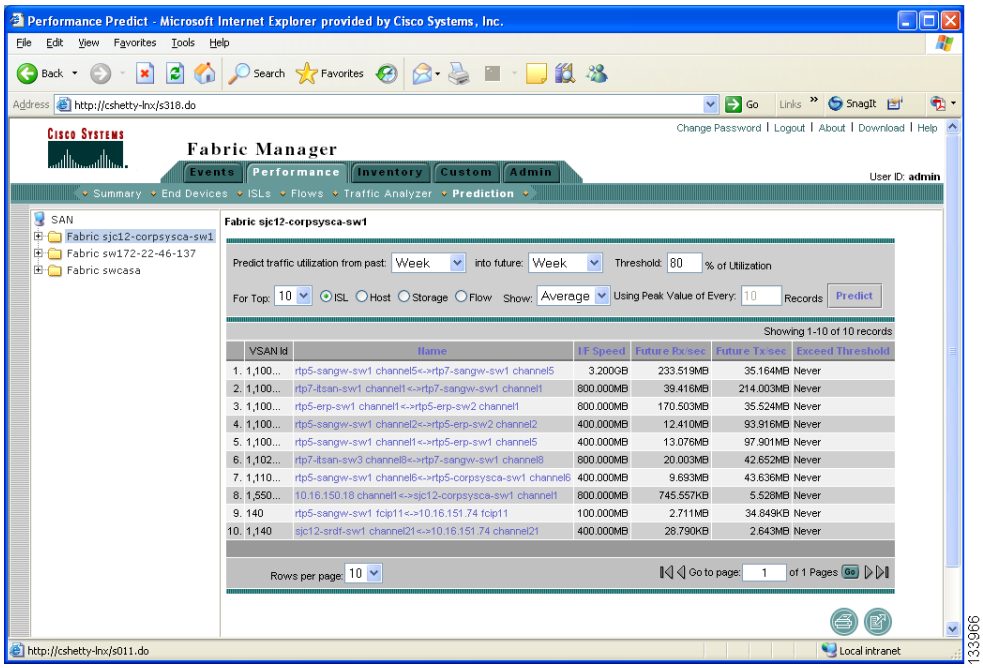

Step 2 Select a value from the **Poll Interval** drop-down list.

Step 3 Decide how you want your data to be interpreted by looking at the **Show Rx/Tx** drop-down menu. The table updates each polling interval to show an overview of the receive and transmit data for each active port on the switch.

Step 4 Select a value from the **Show Rx/Tx** drop-down list. If you select **Util%**, you need to also select values from the two **Show Rx/Tx > %Util/sec** drop-down lists. The first value is the warning level and the second value is the critical threshold level for event reporting.

Note that you can also display percent utilization for a single port by selecting the port and clicking the **Monitor Selected Interface Traffic Util %** icon.

Step 5 Click the **Save Configuration** icon.

The configurable monitor per port option gives statistics for in and out traffic on that port, errors, class 2 traffic and other data that can be graphed over a period of time to give a real-time view into the performance of the port.

To configure per port monitoring using Device Manager, follow these steps:

Step 1 Click the **Device** tab.

You see the graphic representation of the ports shown in Figure 54-2.

Figure 54-2 Device Manager Device Tab

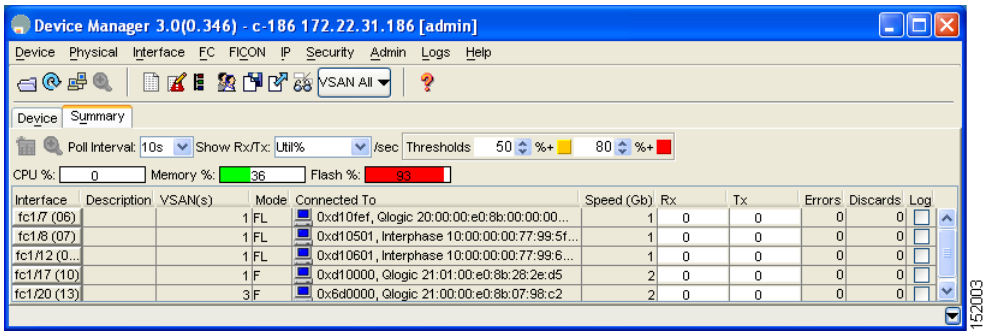

Step 2 Right-click the port you are interested in and choose **Monitor** from the drop-down menu.

You see the port real-time monitor dialog box shown in Figure 54-3.

Figure 54-3 Device Manager Monitor Dialog Box

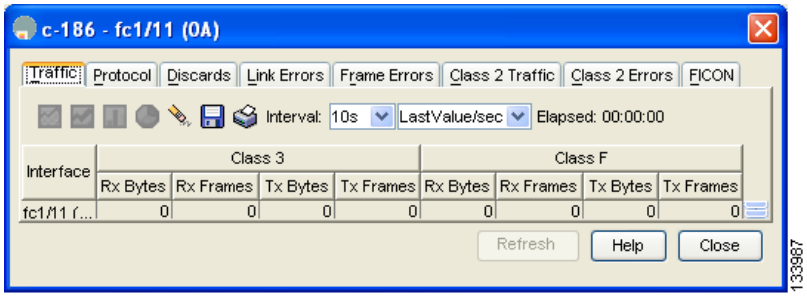

Step 3 Select a value from the **Interval** drop-down list to determine how often data is updated in the table shown here.

Step 4 Click a statistical value in the table then click one of the graphing icons to display a running graph of that statistic over time. You see a graph window that contains options to change the graph type.

Tip You can open multiple graphs for statistics on any of the active ports on the switch.

Fabric Manager Real-Time ISL Statistics

You can configure Fabric Manager to gather ISL statistics in real time. These ISL statistics include receive and transmit utilization, bytes per second, as well as errors and discards per ISL.

To configure ISL statistics using Fabric Manager, follow these steps:

Step 1 Select **Performance > ISLs in Real-Time**.

You see any ISL statistics in the Information pane (see Figure 54-4).

Figure 54-4 ISL Performance in Real Time

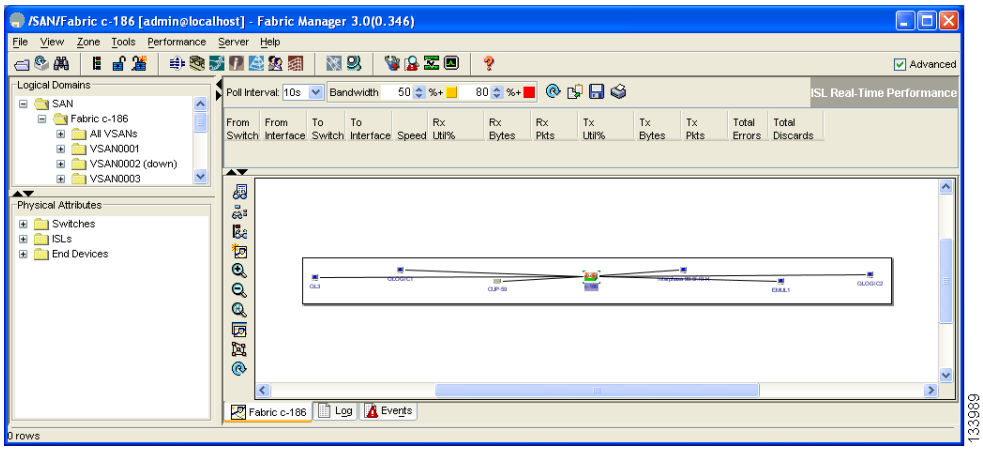

Step 2 Select a value from the **Poll Interval** drop-down list.

Step 3 Select two values from the **Bandwidth** utilization thresholds drop-down lists, one value for the minor threshold and one value for the major threshold.

The table shown updates each polling interval to show the statistics for all configured ISLs in the fabric.

Step 4 Select a row in the table to highlight that ISL in blue in the Topology map.

Historical Performance Monitoring

Performance Manager gathers network device statistics historically and provides this information graphically using a web browser. It presents recent statistics in detail and older statistics in summary. Performance Manager also integrates with external tools such as Cisco Traffic Analyzer.

The Performance Manager has three operational stages:

- Definition—Uses two configuration wizards to create a collection configuration file.
- Collection—Reads the configuration file and collects the desired information.
- Presentation—Generates web pages to present the collected data.

See the "Performance Manager Architecture" section on page 7-1 for an overview of Performance Manager.

Creating a Flow with Performance Manager

Performance Manager has a Flow Configuration Wizard that steps you through the process of creating host-to-storage, storage-to-host, or bidirectional flows. Table 54-1 explains the Flow Type radio button that defines the type of traffic monitored.

Table 54-1 Performance Manager Flow Types

Flow type	Description
Host->Storage	Unidirectional flow, monitoring data from the host to the storage element
Storage->Host	Unidirectional flow, monitoring data from the storage element to the host
Both	Bidirectional flow, monitoring data to and from the host and storage elements.

Once defined, these flows can be added to a collection configuration file to monitor the traffic between a host/storage element pair.

To create a flow using Fabric Manager, follow these steps:

Step 1 Choose **Performance > Create Flows** to launch the Flow Configuration Wizard shown in Figure 54-5.

Figure 54-5 Flow Configuration Wizard

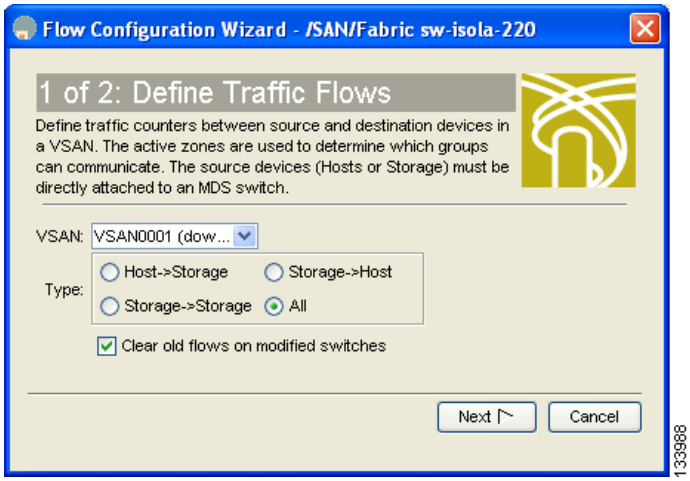

Step 2 Choose the VSAN from which you want to create flows. Flows are defined per VSAN.

Step 3 Choose one of the **Type** radio buttons for the flow type you want to define (**Host->Storage**, **Storage->Host**, **Storage->Storage**, or **All.**

Step 4 Check the **Clear old flows on modified switches** check box if you want to remove old flow data.

Step 5 Click **Next** to review the available flows for the chosen VSAN.

Step 6 Remove any flows you are not interested in.

Step 7 Click **Finish** to create the flow.

The flows created become part of the collection options in the Performance Manager Configuration Wizard.

Creating a Collection with Performance Manager

The Performance Manager Configuration Wizard steps you through the process of creating collections using configuration files. Collections are defined for one or all VSANs in the fabric. Collections can include statistics from the SAN element types described in Table 54-2.

Table 54-2 Performance Manager Collection Types

Collection Type	Description
ISLs	Collects link statistics for ISLs.
Host	Collects link statistics for SAN hosts.
Storage	Collects link statistics for a storage elements.
Flows	Collects flow statistics defined by the Flow Configuration Wizard.

Using Performance Thresholds

The Performance Manager Configuration Wizard allows you to set up two thresholds that trigger events when the monitored traffic exceeds the percent utilization configured. These event triggers can be set as either Critical or Warning events that are reported on the Fabric Manager web client Events browser page.

You must choose either absolute value thresholds or baseline thresholds that apply to all transmit or receive traffic defined in the collection. Click the **Use absolute values** radio button on the last screen of the Performance Manager Configuration Wizard to configure thresholds that apply directly to the statistics gathered. These statistics, as a percent of the total link capacity, are compared to the percent utilization configured for the threshold type. If the statistics exceed either configured threshold, an event is shown on the Fabric Manager web client Events tab.

As an example, the collection has absolute value thresholds set for 60% utilization (for warning) and 80% utilization (for critical). If Performance Manager detects that the traffic on a 1-Gigabit link in its collection exceeds 600 Mbps, a warning event is triggered. If the traffic exceeds 800 Mbps, a critical event is triggered.

Baseline thresholds are defined for a configured time of day or week (1 day, 1 week, or 2 weeks). The baseline is created by calculating the average of the statistical results for the configured time each day, week, or every 2 weeks. Table 54-3 shows an example of the statistics used to create the baseline value for a collection defined at 4 pm on a Wednesday.

Table 54-3 Baseline Time Periods for a Collection Started on Wednesday at 4pm

Baseline Time Window	Statistics Used in Average Calculation
1 day	Every prior day at 4 pm
1 week	Every prior Wednesday at 4 pm
2 weeks	Every other prior Wednesday at 4 pm

Baseline thresholds create a threshold that adapts to the typical traffic pattern for each link for the same time window each day, week, or every 2 weeks. Baseline thresholds are set as a percent of the average (110% to 500%), where 100% equals the calculated average.

As an example, a collection is created at 4 pm on Wednesday, with baseline thresholds set for 1 week, at 150% of the average (warning) and 200% of the average (critical). Performance Manager recalculates the average for each link at 4 pm every Wednesday by taking the statistics gathered at that time each Wednesday since the collection started. Using this as the new average, Performance Manager compares each received traffic statistic against this value and sends a warning or critical event if the traffic on a link exceeds this average by 150% or 200% respectively.

Table 54-4 shows two examples of 1-Gigabit links with different averages in our example collection and at what traffic measurements the Warning and Critical events are sent.

Table 54-4 Example of Events Generated for 1-Gigabit Links

Average	Warning Event Sent at 150%	Critical Event Sent at 200%
400 Mbps	600 Mbps	800 Mbps
200 Mbps	300 Mbps	400 Mbps

Set these thresholds on the last screen of the Collections Configuration Wizard by checking the **Send events if traffic exceeds threshold** check box.

Using the Performance Manager Configuration Wizard

See the "Creating Performance Collections" section on page 6-48.

Viewing Performance Manager Reports

You can view Performance Manager statistical data using preconfigured reports that are built on demand and displayed in a web browser. These reports provide summary information as well as detailed statistics that can be viewed for daily, weekly, monthly, or yearly results.

Choose **Performance > Reports** to access Performance Manager reports from Fabric Manager. This opens a web browser window showing the default Fabric Manager web client event summary report. Click the **Performance** tab to view the Performance Manager reports. Performance Manager begins reporting data ten minutes after the collection is started

Note Fabric Manager Web Services must be running for this to work. See the "Launching Fabric Manager Web Services" section on page 6-5.

Performance Summary

The Performance Summary page presents a dashboard display of the throughput and link utilization for hosts, ISLs, storage, and flows for the last 24-hour period. The summary provides a quick overview of the fabric's bandwidth consumption and highlights any hotspots.

The report includes network throughput pie charts and link utilization pie charts. Use the navigation tree on the left to show summary reports for monitored fabrics or VSANs. The summary displays charts for all hosts, storage elements, ISLs, and flows. Each pie chart shows the percent of entities (links, hosts, storage, ISLs, or flows) that measure throughput or link utilization on each of six predefined ranges. Move the mouse over a pie chart section to see how many entities exhibit that range of statistics. Double-click any pie chart to bring up a table of statistics for those hosts, storage elements, ISLs, or flows.

Performance Tables and Details Graphs

Click **Host, Storage, ISL**, or **Flow** to view traffic over the past day for all hosts, storage, ISLs, or flows respectively. A table lists all of the selected entities, showing transmit and receive traffic and errors and discards, if appropriate. The table can be sorted by any column heading. The table can also be filtered by day, week, month, or year. Tables for each category of statistics display average and peak throughput values and provide hot-links to more detailed information.

Clicking a link in any of the tables opens a details page that shows graphs for traffic by day, week, month, and year. If flows exist for that port, you can see which storage ports sent data. The details page also displays graphs for errors and discards if they are part of the statistics gathered and are not zero.

If you double-click a graph on a Detail report, it will launch the Cisco Traffic Analyzer for Fibre Channel, if configured. The aliases associated with hosts, storage devices, and VSANs in the fabric are passed to the Cisco Traffic Analyzer to provide consistent, easy identification.

Viewing Performance of Host-Optimized Port Groups

You can monitor the performance of host-optimized port groups by clicking **Performance > End Devices** and selecting **Port Groups** from the Type drop-down list.

Viewing Performance Manager Events

Performance Manager events are viewed through Fabric Manager Web Services.

To view problems and events in Fabric Manager Web Services, follow this step:

Step 1 Choose a fabric and then click the **Events** tab to see a summary or detailed report of the problems and events that have occurred in the selected fabric.

Generating Top10 Reports in Performance Manager

Cisco MDS SAN-OS Release 2.1(1a) introduces the ability to generate historical Top10 reports that can be saved for later review. These reports list the entities from the data collection, with the most active entities appearing first. This is a static, one-time only report that generates averages and graphs of the data collection as a snapshot at the time the report is generated.

Tip Name the reports with a timestamp so that you can easily find the report for a given day or week.

These Top10 reports differ from the other monitoring tables and graphs in Performance Manager in that the other data is continuously monitored and is sortable on any table column. The Top10 reports are a snapshot view at the time the report was generated.

Note Top10 reports require analyzing the existing data over an extended period of time and can take hours or more to generate on large fabrics.

To generate a Top10 report using Fabric Manager, follow these steps:

Step 1 Click **Performance > Reports**.

You see the Fabric Manager Web Client login screen.

Step 2 Log in to Fabric Manager Web Services.

Step 3 Click the **Custom** tab and select the template **Top10_Hosts**.

Step 4 Click the **Create** tab.

You see the screen to create a new top ten report shown in Figure 54-6.

Figure 54-6 Create a Top Ten Report

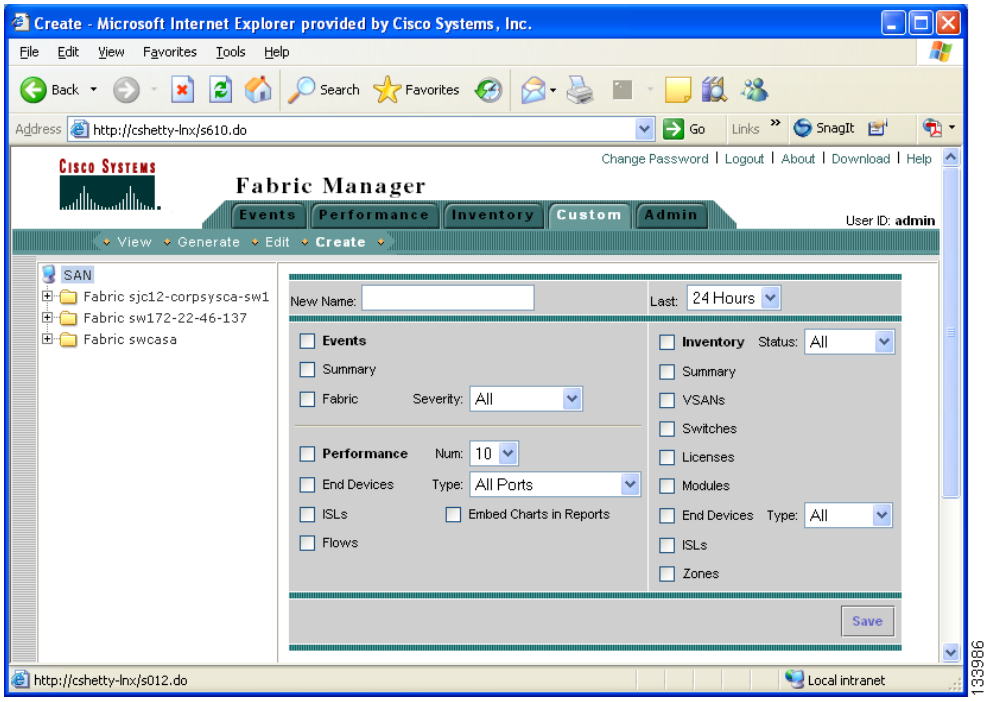

Step 5 Enter a name for a new report.

Step 6 Indicate the details of the report using the check boxes and drop-down menus.

Step 7 Click **Save**.

The report may take hours to generate. When finished, the report appears by name in the left pane navigation bar.

Step 8 Click the name of the generated report to see the Top10 tables for your fabric.

Step 9 Click the name of any entity in the Top10 tables to see a series of graphs for the transmit and receive data rates as well as errors and discards.

Note Fabric Manager Web Services must be running for this to work. See the "Launching Fabric Manager Web Services" section on page 6-5.

Generating Top10 Reports Using Scripts

You can generate Top10 reports manually by issuing the following commands:

- On UNIX, run the script:

```
"/<user_directory>/.cisco_mds9000/bin/pm.sh display pm/pm.xml <output_directory>"
```

- On Windows, run the script:

```
"c:\Program Files\Cisco Systems\MDS 9000\bin\pm.bat display pm\pm.xml
<output_directory>"
```

On UNIX, you can automate the generation of the Top10 reports on your Fabric Manager Server host by adding the following cron entry to generate the reports once an hour:

```
0 * * * * /<user_directory>/.cisco_mds9000/bin/pm.sh display pm/pm.xml <output_directory>
```

If your crontab does not run automatically or Java complains about an exception similar to Example 54-1, you need to add "-Djava.awt.headless=true" to the JVMARGS command in /<user_directory>/.cisco_mds9000/bin/pm.sh.

Example 54-1 Example Java Exception

```
in thread "main" java.lang.InternalError Can't connect to X11 window server using '0.0' as
the value of the DISPLAY variable.
```

Exporting Data Collections to XML Files

The RRD files used by Performance Manager can be exported to a freeware tool called rrdtool. The rrd files are located in pm/db on the Fabric Manager Server. To export the collection to an XML file, enter the following command at the operating system command-line prompt:

```
/bin/pm.bat xport xxx yyy
```

In this command, *xxx* is the RRD file and *yyy* is the XML file that is generated. This XML file is in a format that rrdtool is capable of reading with the command:

```
rrdtool restore filename.xml filename.rrd
```

You can import an XML file with the command:

```
bin/pm.bat pm restore <xmlFile> <rrdFile>
```

This reads the XML export format that rrdtool is capable of writing with the command:

```
rrdtool xport filename.xml filename.rrd.
```

The **pm xport** and **pm restore** commands can be found on your Fabric Manager Server at bin\PM.bat for Windows platforms or bin/PM.sh on UNIX platforms. For more information on the rrdtool, refer to the following website: http://www.rrdtool.org.

Exporting Data Collections in Readable Format

Cisco MDS SAN-OS Release 2.1(1a) introduces the ability to export data collections in comma-separated format (CSV). This format can be imported to various tools, including Microsoft Excel. You can export these readable data collections either from the Fabric Manager Web Services menus or in batch mode from the command line on Windows or UNIX. Using Fabric Manager Web Services, you can export one file. Using batch mode, you can export all collections in the pm.xml file.

Note Fabric Manager Web Services must be running for this to work. See the "Launching Fabric Manager Web Services" section on page 6-5.

To export data collections to Microsoft Excel using Fabric Manager Web Services, follow these steps:

Step 1 Click the **Performance** tab on the main page.

You see the overview table.

Step 2 Click the **Flows** sub-tab.

Step 3 Right-click the name of the entity you want to export and select **Export to Microsoft Excel**.

You see the Excel chart for that entity in a pop-up window.

To export data collections using command-line batch mode, follow these steps:

Step 1 Go to the installation directory on your workstation and then go to the bin directory.

Step 2 On Windows, enter **.\pm.bat export C:\Program Files\Cisco Systems\MDS 9000\pm\pm.xml <*export directory*>.** This creates the csv file (export.csv) in the *export directory* on your workstation.

Step 3 On UNIX, enter **./pm.sh export /usr/local/cisco_mds9000/pm/pm.xml <*export directory*>**. This creates the csv file (export.csv) in the *export directory* on your workstation.

When you open this exported file in Microsoft Excel, the following information displays:

- Title of the entity you exported and the address of the switch the information came from.
- The maximum speed seen on the link to or from this entity.
- The VSAN ID and maximum speed.
- The timestamp, followed by the receive and transmit data rates in bytes per second.

Configuring Performance Manager for Use with Cisco Traffic Analyzer

Performance Manager works in conjunction with the Cisco Traffic Analyzer to allow you to monitor and manage the traffic on your fabric. Using Cisco Traffic Analyzer with Performance Manager requires the following components:

- A configured Fibre Channel Switched Port Analyzer (SPAN) destination (SD) port to forward Fibre Channel traffic.
- A Port Analyzer Adapter 2 (PAA-2) to convert the Fibre Channel traffic to Ethernet traffic.
- Cisco Traffic Analyzer software to analyze the traffic from the PAA-2.

To configure Performance Manager to work with the Cisco Traffic Analyzer, follow these steps:

Step 1 Set up the Cisco Traffic Analyzer according to the instructions in the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note.*

Step 2 Get the following three items of information:

- The IP address of the management workstation on which you are running Performance Manager and Cisco Traffic Analyzer.
- The path to the directory where Cisco Traffic Analyzer is installed.
- The port that is used by Cisco Traffic Analyzer (the default is 3000).

Step 3 Start the Cisco Traffic Analyzer.

a. Choose **Performance > Traffic Analyzer > Open**.

b. Enter the URL for the Cisco Traffic Analyzer, in the format

`http://<ip address>:<port number>`

where:

ip address is the address of the management workstation on which you have installed the Cisco Traffic Analyzer, and

:*port number* is the port that is used by Cisco Traffic Analyzer (the default is :3000).

c. Click **OK**.

d. Choose **Performance > Traffic Analyzer > Start**.

e. Enter the location of the Cisco Traffic Analyzer, in the format

`D:\<directory>\ntop.bat`

where:

D: is the drive letter for the disk drive where the Cisco Traffic Analyzer is installed, and

directory is the directory containing the ntop.bat file.

f. Click **OK**.

Step 4 Create the flows you want Performance Manager to monitor, using the Flow Configuration Wizard. See the "Creating a Flow with Performance Manager" section on page 54-4

Step 5 Define the data collection you want Performance Manager to gather, using the Performance Manager Configuration Wizard. See the "Creating a Collection with Performance Manager" section on page 54-6.

a. Choose the VSAN you want to collect information for or choose **All VSANs**.

b. Check the types of items you want to collect information for (Hosts, ISLs, Storage Devices, and Flows).

c. Enter the URL for the Cisco Traffic Analyzer in the format

```
http://<ip address>/<directory>
```

where:

ip address is the address of the management workstation on which you have installed the Cisco Traffic Analyzer, and *directory* is the path to the directory where the Cisco Traffic Analyzer is installed.

d. Click **Next**.

e. Review the data collection on this and the next section to make sure this is the data you want to collect.

f. Click **Finish** to begin collecting data.

Note Data is not collected for JBOD or for virtual ports. If you change the data collection configuration parameters during a data collection, you must stop and restart the collection process for your changes to take effect.

Step 6 Choose **Performance > Reports** to generate a report. Performance Manager Web Services must be running. See the "Launching Fabric Manager Web Services" section on page 6-5. You see Web Services; click **Custom** then select a report template.

Note It takes at least five minutes to start collecting data for a report. Do not attempt to generate a report in Performance Manager during the first five minutes of collection.

Step 7 Click the **Cisco Traffic Analyzer** at the top of the Host or Storage detail pages to view the Cisco Traffic Analyzer information, or choose **Performance > Traffic Analyzer > Open**. The Cisco Traffic Analyzer page will not open unless ntop has been started already.

Note For information on capturing a SPAN session and starting a Cisco Traffic Analyzer session to view it, refer to the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note.*

Note For information on viewing and interpreting your Performance Manager data, see the "Historical Performance Monitoring" section on page 54-4.

For information on viewing and interpreting your Cisco Traffic Analyzer data, refer to the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*.

For performance drill-down, Fabric Manager Server can launch the Cisco Traffic Analyzer in-context from the Performance Manager graphs. The aliases associated with hosts, storage devices, and VSANs are passed to the Cisco Traffic Analyzer to provide consistent, easy identification.

CHAPTER 55

Configuring RMON

RMON is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. You can use the RMON alarms and events to monitor Cisco MDS 9000 Family switches running the Cisco SAN-OS Release 2.0(1b) or later software.

This chapter includes the following sections:

About RMON

All switches in the Cisco MDS 9000 Family support the following RMON functions (defined in RFC 2819):

- Alarm—Monitors a specific management information base (MIB) object for a specified interval, triggers an alarm at a specified value (rising threshold), and resets the alarm at another value (falling threshold). Alarms can be used with events; the alarm triggers an event, which can generate a log entry or an SNMP trap.
- Event—Determines the action to take when an event is triggered by an alarm. The action can be to generate a log entry, an SNMP trap, or both.

Refer to the *Cisco MDS 9000 Family MIB Quick Reference* for agent and management information.

See the "About SNMP" section on page 34-2 for SNMP security-related CLI configurations.

Configuring RMON Using Threshold Manager

RMON is disabled by default and no events or alarms are configured in the switch. You can configure your RMON alarms and events by using the CLI or by using Threshold Manager in Device Manager.

The Threshold Monitor allows you to trigger an SNMP event or log a message when the selected statistic goes over a configured threshold value. RMON calls this a rising alarm threshold. The configurable settings are:

- Variable—The statistic you want to set the threshold value on.

- Value—The value of the variable that you want the alarm to trigger at. This value is the difference (delta) between two consecutive polls of the variable by Device Manager.
- Sample—The sample period (in seconds) between two consecutive polls of the variable. Select your sample period such that the variable would not cross the threshold value you set under normal operating conditions.
- Warning—The warning level used by Device Manager to indicate the severity of the triggered alarm. This is a Fabric Manager and Device Manager enhancement to RMON.

Note To configure any type of RMON alarm (absolute or delta, rising or falling threshold) click **More** on the Threshold Manager dialog box. You should be familiar with how RMON defines these concepts before configuring these advanced alarm types. Refer to the RMON-MIB (RFC 2819) for information on how to configure RMON alarms.

You must also configure SNMP on the switch to access RMON MIB objects.

RMON Alarm Configuration

Threshold Manager provides a list of common MIB objects that you may want to set an RMON threshold and alarm on. You can also set an alarm on any MIB object. The specified MIB must be an existing SNMP MIB object in standard dot notation (1.3.6.1.2.1.2.2.1.14.16 for ifInOctets.16).

Use one of the following options to specify the interval to monitor the MIB variable (ranges from 1 to 4294967295 seconds):

- Use the **delta** option to test the change between samples of a MIB variable.
- Use the **absolute** option to test each MIB variable directly.
- Use the **delta** option to test any MIB objects that are counters.

The range for the **rising threshold** and **falling threshold** values is -2147483647 to 2147483647.

Caution The **falling threshold** must be less than the **rising threshold**.

You can optionally specify the following parameters:

- The event-number to trigger if the rising or falling threshold exceeds the specified limit.
- The owner of the alarm.

Enabling RMON Alarms by Port

To configure an RMON alarm for one or more ports using Device Manager, follow these steps:

Step 1 Choose **Admin > Events > Threshold Manager** and click the **FC Interfaces** tab.

You see the Threshold Manager dialog box shown in Figure 55-1.

Figure 55-1 Threshold Manager Dialog Box

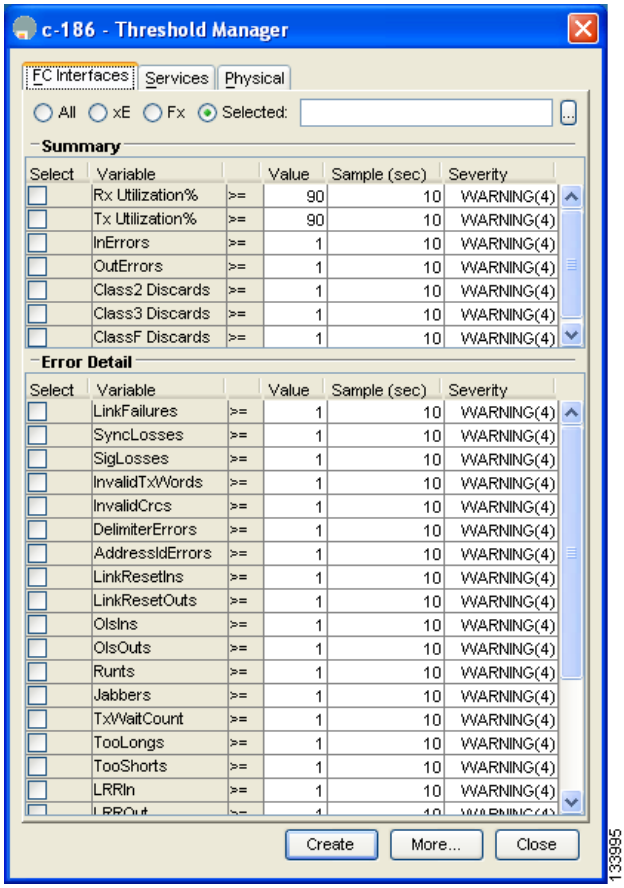

Step 2 Choose the **Selected** radio button to select individual ports for this threshold alarm (see Figure 55-1).

a. Click the **...** button to the right of the Selected field to display all ports.

b. Select the ports you want to monitor.

c. Click **OK** to accept the selection.

Alternatively, click the appropriate radio button to choose ports by type: **All** ports, **xE** ports, or **Fx** ports.

Step 3 Check the check box for each variable to be monitored.

Step 4 Enter the threshold value in the Value column.

Step 5 Enter the sampling period in seconds. This is the time between each snapshot of the variable.

Step 6 Choose one of the following severity levels to assign to the alarm: **Fatal**, **Warning**, **Critical**, **Error**, **Information,**

Step 7 Click **Create,**

Step 8 Confirm the operation to define an alarm and a log event when the system prompts you to define a severity event.If you do not confirm the operation, the system only defines a log event.

Step 9 Click **More** then click the **Alarms** tab from the Threshold Manager dialog box to verify the alarm you created.

Step 10 Close both dialog pop-ups.

Enabling RMON Alarms for VSANs

To enable an RMON alarm for one or more VSANs using Device Manager, follow these steps:

Step 1 Choose **Admin > Events > Threshold Manager** and click the **Services** tab.

You see the Threshold Manager dialog box.

Step 2 Click the **Services** tab.

You see the dialog box shown in Figure 55-2.

Figure 55-2 Threshold Manager Services Tab

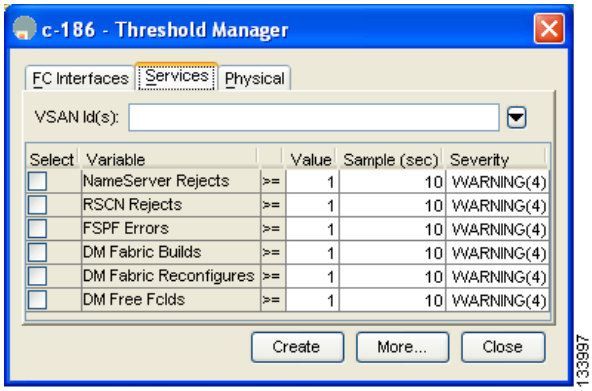

Step 3 Enter one or more VSANs (multiple VSANs separated by commas) to monitor in the VSAN ID(s) field. Use the down arrow to see a list of available VSANs to choose from.

Step 4 Check the check box in the Select column for each variable to monitor.

Step 5 Enter the threshold value in the Value column.

Step 6 Enter the sampling period in seconds.

Step 7 Choose a severity level to assign to the alarm (**FATAL**, **CRITICAL**, **ERROR**, **WARNING**, or **INFORMATION**).

Step 8 Click **Create.**

Step 9 Confirm the operation to define an alarm and a log event when the system prompts you to define a severity event.

If you do not confirm the operation, the system only defines a log event.

Step 10 Click **More** then click the **Alarms** tab in the Threshold Manager dialog box to verify the alarm you created. See Figure 55-3.

Figure 55-3 List of Threshold Alarms

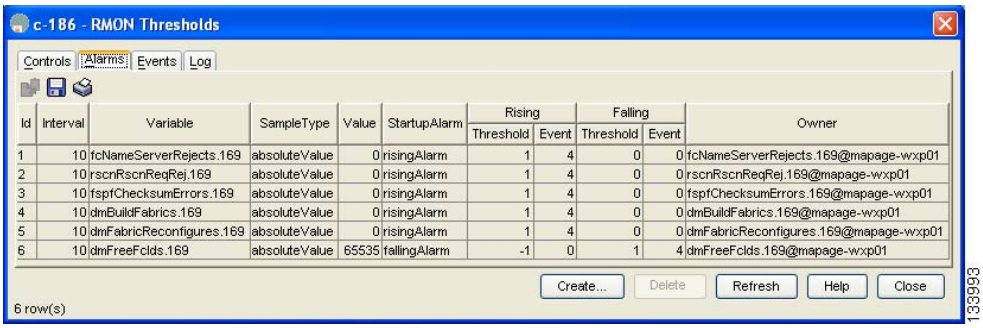

Step 11 Close both pop-up windows.

Enabling RMON Alarms for Physical Components

To configure an RMON alarm for a physical component using Device Manager, follow these steps:

Step 1 Choose **Admin > Events > Threshold Manager** and click the **Physical** tab.

You see the Threshold Manager dialog box with the Physical tab selected (see Figure 55-4).

Figure 55-4 Threshold Manager Physical Tab

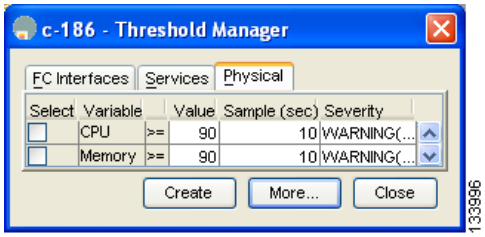

Step 2 Check the check box in the Select column for each variable to monitor.

Step 3 Enter the threshold value in the Value column.

Step 4 Enter the sampling period in seconds.

Step 5 Choose one of the following severity levels to assign to the alarm: FATAL(1), WARNING(2), CRITICAL(3), ERROR(4), INFORMATION(5).

Step 6 Click **Create**.

Step 7 Confirm the operation to define an alarm and a log event when the system prompts you to define a severity event.

If you do not confirm the operation, the system only defines a log event.

Step 8 Click **More** and select the **Alarms** tab from the Threshold Manager dialog box to verify the alarm you created (see Figure 55-5).

Figure 55-5 List of Threshold Alarms

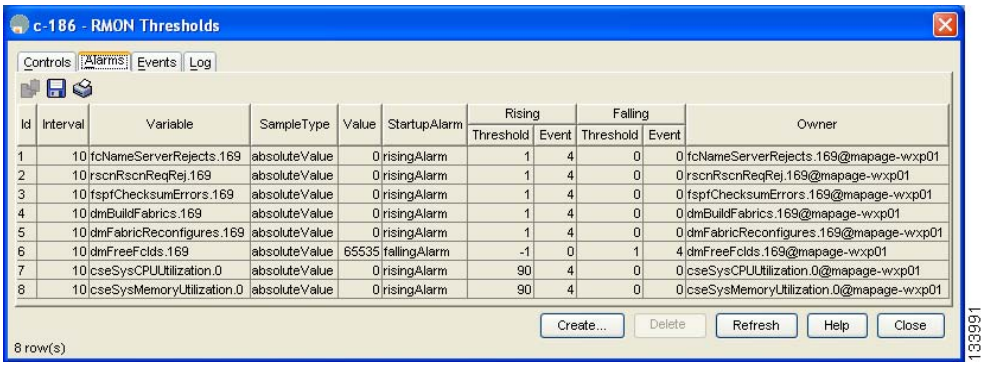

Step 9 Close both dialog boxes.

Managing RMON Events

To define customized RMON events using Device Manager, follow these steps:

Step 1 Choose **Admin > Events > Threshold Manager** and click **More** on the Threshold Manager dialog box.

Step 2 Click the **Events** tab on the RMON Thresholds dialog box.

You see the RMON Events shown in Figure 55-6.

Figure 55-6 RMON Events Thresholds

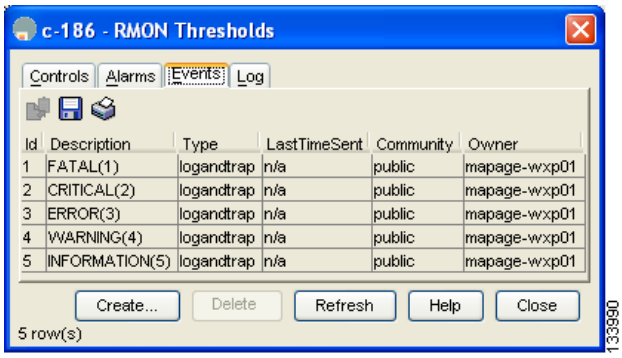

Step 3 Click **Create** to create an event entry.

You see the Create RMON Threshold Event Entry dialog box.

Figure 55-7 Create an RMON Threshold Event

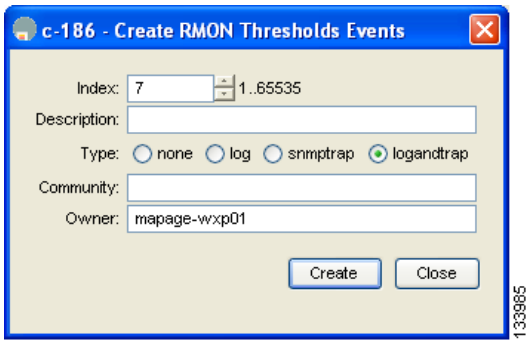

Step 4 Configure the RMON threshold event attributes by choosing the type of event (**log**, **snmptrap**, or **logandtrap**).

Step 5 Increment the index.. If you try to create an event with the existing index, you see a duplicate entry error message.

Step 6 Optionally provide a description and a community.

Step 7 Click **Create** then close this dialog box.

Step 8 Verify that your event is listed in the remaining RMON Thresholds dialog box.

Step 9 Close the RMON Thresholds dialog box

Managing RMON Alarms

To view the alarms that have already been enabled using Device Manager, follow these steps:

Step 1 Choose **Admin > Events > Threshold Manager** and click **More** on the Threshold Manager dialog box.

Step 2 Click the **Alarms** tab.

You see the RMON Thresholds dialog box shown in Figure 55-8.

Figure 55-8 RMON Thresholds Dialog Box 3

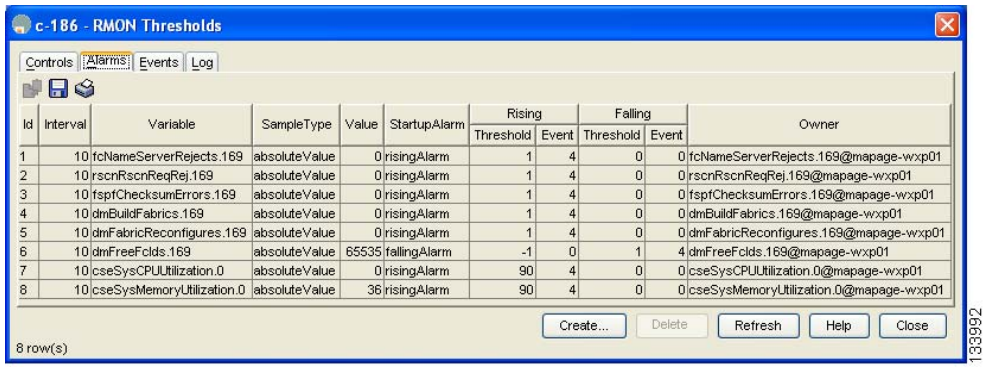

Step 3 Delete any alarm by selecting it, then clicking **Delete**.

Viewing the RMON Log

To view the RMON log using Device Manager, follow these steps:

Step 1 Choose **Admin > Events > Threshold Manager** and click **More** on the Threshold Manager dialog box.

Step 2 Click the **Log** tab in the RMON Thresholds dialog box.

You see the RMON Log dialog box shown in Figure 55-9. This is the log of RMON events that have been triggered by the Threshold Manager.

Figure 55-9 RMON Threshold Manager Log Dialog Box

Default Settings

Table 55-1 lists the default settings for all RMON features in any switch.

Table 55-1 Default RMON Settings

Parameters	Default
RMON alarms	Disabled.
RMON events	Disabled.

CHAPTER 56

Monitoring Network Traffic Using SPAN

This chapter describes the Switched Port Analyzer (SPAN) features provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

About SPAN

The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic though a Fibre Channel interface. Traffic through any Fibre Channel interface can be replicated to a special port called the SPAN destination port (SD port). Any Fibre Channel port in a switch can be configured as an SD port. Once an interface is in SD port mode, it cannot be used for normal data traffic. You can attach a Fibre Channel Analyzer to the SD port to monitor SPAN traffic (see "Configuring the Cisco Fabric Analyzer" section on page 62-18).

SD ports do not receive frames, they merely transmit a copy of the SPAN source traffic. The SPAN feature is non-intrusive and does not affect switching of network traffic for any SPAN source ports (see Figure 56-1).

Figure 56-1 SPAN Transmission

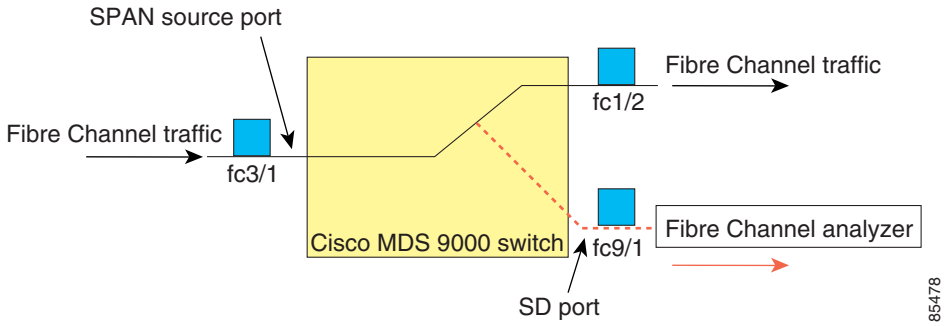

SPAN Sources

SPAN sources refer to the interfaces from which traffic can be monitored. You can also specify VSAN as a SPAN source, in which case, all supported interfaces in the specified VSAN are included as SPAN sources. You can choose the SPAN traffic in the ingress direction, the egress direction, or both directions for any source interface:

- Ingress source (Rx)—Traffic entering the switch fabric through this source interface is *spanned* or copied to the SD port (see Figure 56-2).

Figure 56-2 SPAN Traffic from the Ingress Direction

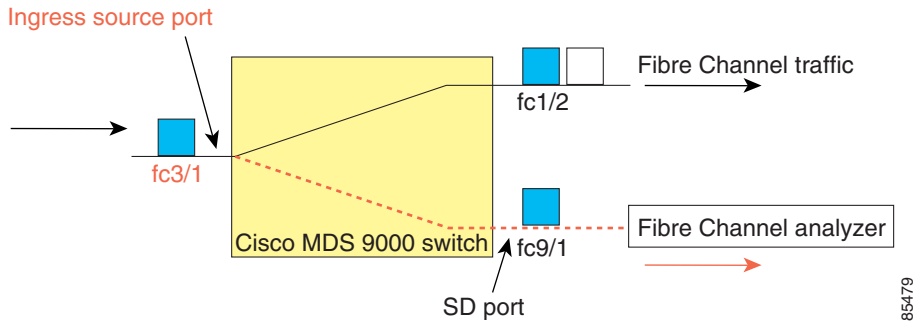

- Egress source (Tx)—Traffic exiting the switch fabric through this source interface is spanned or copied to the SD port (see Figure 56-3).

Figure 56-3 SPAN Traffic from Egress Direction

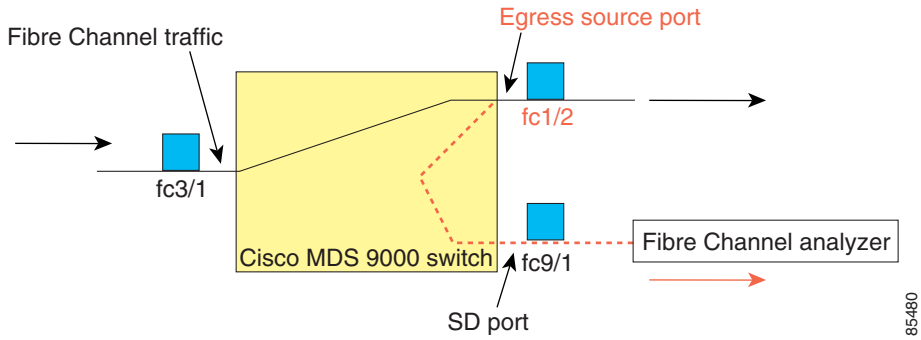

IPS Source Ports

SPAN capabilities are available on the IP Storage Services (IPS) module. The SPAN feature is only implemented on the FCIP and iSCSI virtual Fibre Channel port interfaces, not the physical Gigabit Ethernet ports. You can configure SPAN for ingress traffic, egress traffic, or traffic in both directions for all eight iSCSI and 24 FCIP interfaces that are available in the IPS module.

Note You can configure SPAN for Ethernet traffic using Cisco switches or routers connected to the Cisco MDS 9000 Family IPS modules.

CSM Source Ports

SPAN capabilities are available on the Caching Services Module (CSM).

Refer to the *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for more information.

Allowed Source Interface Types

The SPAN feature is available for the following interface types:

- Physical ports such as F ports, FL ports, TE ports, E ports, and TL ports.
- Interface sup-fc0 (traffic to and from the supervisor):
  - The Fibre Channel traffic from the supervisor module to the switch fabric through the sup-fc0 interface is called ingress traffic. It is spanned when sup-fc0 is chosen as an ingress source port.
  - The Fibre Channel traffic from the switch fabric to the supervisor module through the sup-fc0 interface is called egress traffic. It is spanned when sup-fc0 is chosen as an egress source port.
- PortChannels
  - All ports in the PortChannel are included and spanned as sources.
  - You cannot specify individual ports in a PortChannel as SPAN sources. Previously configured SPAN-specific interface information is discarded.
- IPS module specific Fibre Channel interfaces.
  - iSCSI interfaces
  - FCIP interfaces

VSAN as a Source

When a VSAN as a source is specified, then all physical ports and PortChannels in that VSAN are included as SPAN sources. A TE port is included only when the port VSAN of the TE port matches the source VSAN. A TE port is excluded even if the configured allowed VSAN list may have the source VSAN, but the port VSAN is different.

You cannot configure source interfaces (physical interfaces, PortChannels, or sup-fc interfaces) and source VSANs in the same SPAN session.

Guidelines to Configure VSANs as a Source

The following guidelines apply when configuring VSANs as a source:

- Traffic on all interfaces included in a source VSAN is spanned only in the ingress direction.
- If a VSAN is specified as a source, you cannot perform interface-level SPAN configuration on the interfaces that are included in the VSAN. Previously configured SPAN-specific interface information is discarded.
- If an interface in a VSAN is configured as a source, you cannot configure that VSAN as a source. You must first remove the existing SPAN configurations on such interfaces before configuring VSAN as a source.

- Interfaces are only included as sources when the port VSAN matches the source VSAN. Figure 56-4 displays a configuration using VSAN 2 as a source:
  - All ports in the switch are in VSAN 1 except fc1/1.
  - Interface fc1/1 is the TE port with port VSAN 2. VSANs 1, 2, and 3 are configured in the allowed list.
  - VSAN 1 and VSAN 2 are configured as SPAN sources.

Figure 56-4 VSAN as a Source

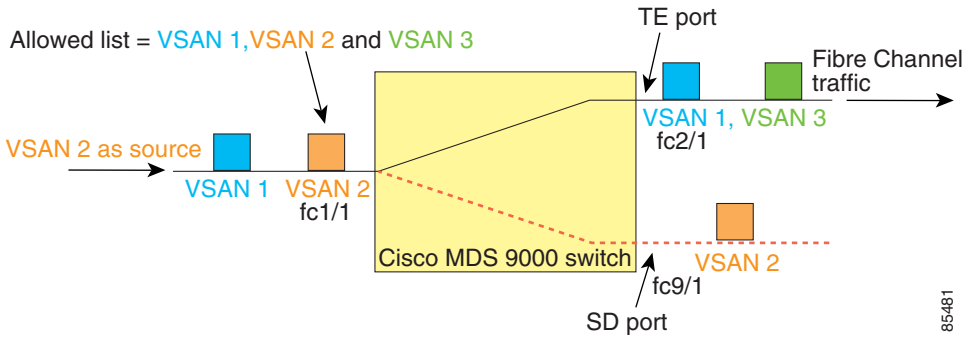

For this configuration, the following apply:

- VSAN 2 as a source includes only the TE port fc1/1 that has port VSAN 2.
- VSAN 1 as a source does not include the TE port fc1/1 as the port VSAN does not match VSAN 1.

  See the "Configuring an Allowed-Active List of VSANs" section on page 20-7 or the "Creating VSANs Statically" section on page 23-6.

SPAN Sessions

Each SPAN session represents an association of one destination with a set of source(s) along with various other parameters that you specify to monitor the network traffic. One destination can be used by one or more SPAN sessions. You can configure up to 16 SPAN sessions in a switch. Each session can have several source ports and one destination port.

To activate any SPAN session, at least one source and the SD port must be up and functioning. Otherwise, traffic is not directed to the SD port.

Tip A source can be shared by two sessions, however, each session must be in a different direction—one ingress and one egress.

You can temporarily deactivate (suspend) any SPAN session. The traffic monitoring is stopped during this time.

Specifying Filters

You can perform VSAN-based filtering to selectively monitor network traffic on specified VSANs. You can apply this VSAN filter to all sources in a session (see Figure 56-4). Only VSANs present in the filter are spanned.

You can specify session VSAN filters that are applied to all sources in the specified session. These filters are bidirectional and apply to all sources configured in the session.

Guidelines to Specifying Filters

The following guidelines apply to SPAN filters:

- PortChannel configurations are applied to all ports in the PortChannel.
- If no filters are specified, the traffic from all active VSANs for that interface is spanned by default.
- While you can specify arbitrary VSAN filters in a session, traffic can only be monitored on the port VSAN or on allowed-active VSANs in that interface.

SD Port Characteristics

An SD port has the following characteristics:

- Ignores BB_credits.
- Allows data traffic only in the egress (Tx) direction.
- Does not require a device or an analyzer to be physically connected.
- Supports only 1 Gbps or 2 Gbps speeds. The auto speed option is not allowed.
- Multiple sessions can share the same destination ports.
- If the SD port is shut down, all shared sessions stop generating SPAN traffic.
- The outgoing frames can be encapsulated in Extended Inter-Switch Link (EISL) format.
- The SD port does not have a port VSAN.
- SD ports cannot be configured using Storage Services Modules (SSMs).
- The port mode cannot be changed if it is being used for a SPAN session.

Note If you need to change an SD port mode to another port mode, first remove the SD port from all sessions and then change the port mode.

Guidelines to Configure SPAN

The following guidelines apply for SPAN configurations:

- You can configure up to 16 SPAN sessions with multiple ingress (Rx) sources.
- You can configure a maximum of three SPAN sessions with one egress (Tx) port.
- In a 32-port switching module, you must configure the same session in all four ports in one port group (unit). If you wish, you can also configure only two or three ports in this unit.

- SPAN frames are dropped if the sum of the bandwidth of the sources exceeds the speed of the destination port.
- Frames dropped by a source port are not spanned.

Configuring SPAN

To monitor network traffic using SD ports, follow these steps:

Step 1 Configure the SD port.

Step 2 Attach the SD port to a specific SPAN session.

Step 3 Monitor network traffic by adding source interfaces to the session.

To configure an SD port for SPAN monitoring using Device Manager, follow these steps:

Step 1 Right-click the port you want to configure and select **Configure**.

You see the general port configuration dialog.

Step 2 Under Mode, choose **SD**.

Step 3 Click **Apply** to accept the change.

Step 4 Close the dialog box.

Creating SPAN Sessions

To create a SPAN session using Device Manager, follow these steps:

Step 1 Choose **Interface > SPAN**. You see the SPAN dialog box.

Step 2 Click the **Sessions** tab.

Step 3 Click **Create**.

You see the Create SPAN Session dialog box.

Step 4 Choose the session ID (from 1-16) using the up or down arrows, and click **Create**.

The session appears in the SPAN dialog box shown in Figure 56-5.

Figure 56-5 Create a SPAN Session

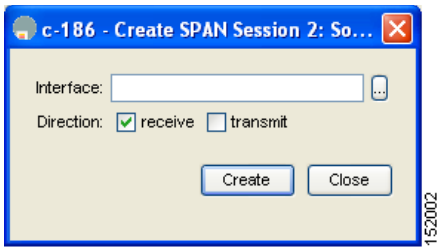

Step 5 Repeat Step 4 for each session you want to create.

Step 6 Choose the destination interface by clicking once in the Dest Interface field for the appropriate session.

Step 7 Choose the filter VSAN list by clicking once in the Filter VSAN List field for the appropriate session.

Step 8 Choose **active** or in **active** admin status in the Admin drop-down list.

Step 9 Click **Apply** to save your changes, or click **Close** to close the SPAN Sessions dialog box without saving your changes.

Step 10 Close the two dialog boxes.

Editing SPAN Sources

To edit a SPAN source using Device Manager, follow these steps:

Step 1 Choose **Interface > SPAN**.

You see the SPAN dialog box.

Step 2 Click the **Sources** tab.

You see the dialog box shown in Figure 56-6.

Figure 56-6 Edit SPAN Sources Dialog Box

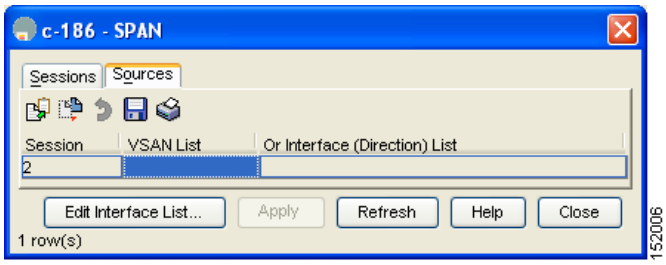

Step 3 Click once in the VSAN List field, and enter the VSAN list name.

Step 4 Click **Edit Interface List**.

You see the Source Interfaces dialog box shown in Figure 56-7.

Figure 56-7 Source Interfaces Dialog Box

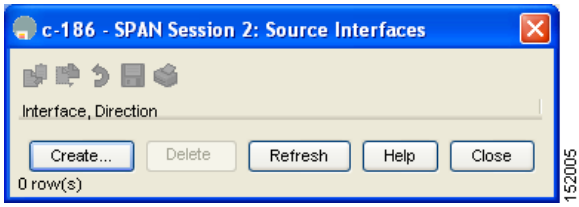

Step 5 Click **Create**.

You see the Create FC Interface Source dialog box.

Step 6 Click the **...** button to display the list of available FC ports.

Step 7 Choose a port and click **OK**.

Step 8 Click the direction (**receive** or **transmit**) you want.

Step 9 Click **Create** to create the FC interface source, or click **Close** to close the Create FC Interface Source dialog box without creating the interface source.

Step 10 Click **Close** in each of the three open dialog boxes.

Deleting SPAN Sessions

To delete a SPAN session using Device Manager, follow these steps.

Step 1 Choose **Interface > SPAN**.

You see the SPAN dialog box.

Step 2 Click the **Sessions** tab.

Step 3 Click the SPAN session you want to delete.

Step 4 Click **Delete**.

The SPAN session is deleted.

Step 5 Close the dialog box.

SPAN Conversion Behavior

As of Cisco MDS SAN-OS Release 1.1(1), SPAN features (configured in any prior release) are converted as follows:

- If source interfaces and source VSANs are configured in a given session, then all the source VSANs are removed from that session.

  For example,

Before Cisco MDS SAN-OS Release 1.0(4):

```
Session 1 (active)
   Destination is fc1/9
   No session filters configured
   Ingress (rx) sources are
      vsans 10-11
     fc1/3,
   Egress (tx) sources are
     fc1/3,
```

Once upgraded to Cisco MDS SAN-OS Release 1.1(1):

```
Session 1 (active)
   Destination is fc1/9
   No session filters configured
   Ingress (rx) sources are
     fc1/3,
   Egress (tx) sources are
     fc1/3,
```

Session 1 had both source interfaces and source VSANs before the upgrade. After the upgrade, the source VSANs were removed (rule 1).

- If interface level VSAN filters are configured in source interfaces, then the source interfaces are also removed from the session. If this interface is configured in both directions, it is removed from both directions.

  For example,

  Before Cisco MDS SAN-OS Release 1.0(4):

```
Session 2 (active)
   Destination is fc1/9
   No session filters configured
   Ingress (rx) sources are
      vsans 12
     fc1/6 (vsan 1-20),
   Egress (tx) sources are
     fc1/6 (vsan 1-20),
```

  Once upgraded to Cisco MDS SAN-OS Release 1.1(1):

```
Session 2 (inactive as no active sources)
   Destination is fc1/9
   No session filters configured
   No ingress (rx) sources
   No egress (tx) sources
```

Note The deprecated configurations are removed from persistent memory once a switchover or a new startup configuration is implemented.

  Session 2 had a source VSAN 12 and a source interface fc1/6 with VSAN filters specified in Cisco MDS SAN-OS Release 1.0(4). When upgraded to Cisco MDS SAN-OS Release 1.1(1) the following changes are made:

  - The source VSAN (VSAN 12) is removed (rule 1).
  - The source interface fc1/6 had VSAN filters specified—it is also removed (rule 2).

Monitoring Traffic Using Fibre Channel Analyzers

You can use SPAN to monitor traffic on an interface without any traffic disruption. This feature is specially useful in troubleshooting scenarios where traffic disruption changes the problem environment and makes it difficult to reproduce the problem.

Without SPAN

You can monitor traffic using interface fc1/1 in a Cisco MDS 9000 Family switch that is connected to another switch or host. You need to physically connect a Fibre Channel analyzer between the switch and the storage device to analyze the traffic through interface fc1/1 as shown in Figure 56-8.

Figure 56-8 Fibre Channel Analyzer Usage Without SPAN

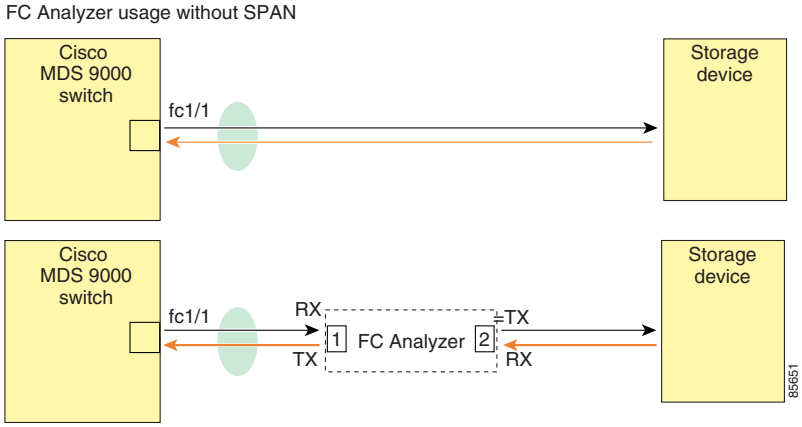

This type of connection has the following limitations:

- It requires you to physically insert the FC analyzer between the two network devices.
- It disrupts traffic when the Fibre Channel analyzer is physically connected.
- The analyzer captures data only on the Rx links in both port 1 and port 2. Port 1 captures traffic exiting interface fc1/1 and port 2 captures ingress traffic into interface fc1/1.

With SPAN

Using SPAN you can capture the same traffic scenario shown in Figure 56-8 without any traffic disruption. The Fibre Channel analyzer uses the ingress (Rx) link at port 1 to capture all the frames going out of the interface fc1/1. It uses the ingress link at port 2 to capture all the ingress traffic on interface fc1/1.

Using SPAN you can monitor ingress traffic on fc1/1 at SD port fc2/2 and egress traffic on SD port fc2/1. This traffic is seamlessly captured by the FC analyzer as shown in Figure 56-9.

Figure 56-9 Fibre Channel Analyzer Using SPAN

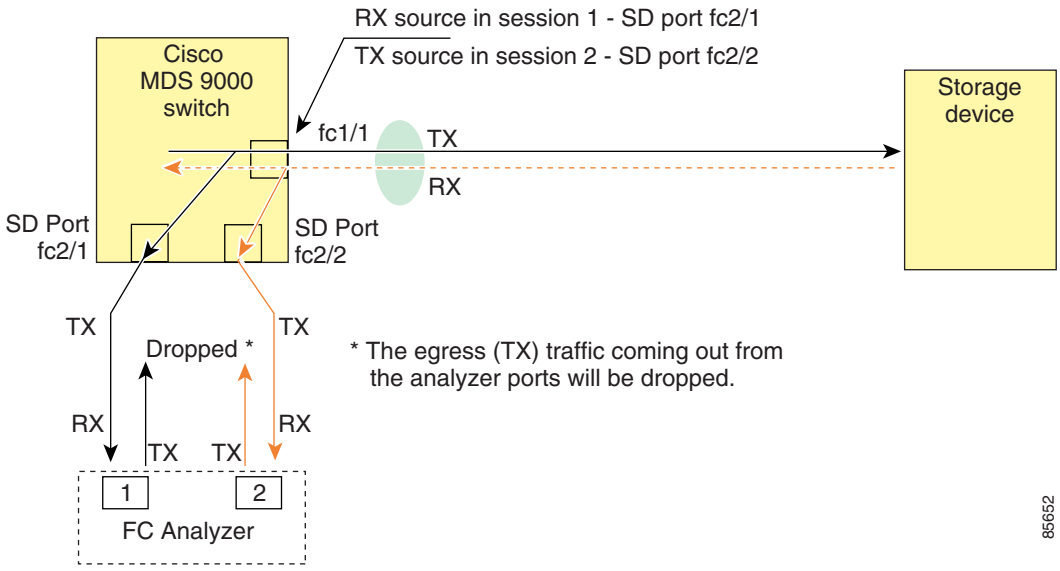

Configuring Analyzers Using SPAN

To configure Fibre Channel Analyzers using SPAN for the example in Figure 56-9, follow these steps:

Step 1 Configure SPAN on interface fc1/1 in the ingress (Rx) direction to send traffic on SD port fc2/1 using session 1.

Step 2 Configure SPAN on interface fc1/1in the egress (Tx) direction to send traffic on SD port fc2/2 using session 2.

Step 3 Physically connect fc2/1 to port 1 on the Fibre Channel analyzer.

Step 4 Physically connect fc2/2 to port 2 on the Fibre Channel analyzer.

Single SD Port to Monitor Traffic

You do not need to use two SD ports to monitor bidirectional traffic on any interface as shown in Figure 56-9. You can use one SD port and one FC analyzer port by monitoring traffic on the interface at the same SD port fc2/1.

Figure 56-10 shows a SPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress directions. This setup is more advantageous and cost effective than the setup shown in Figure 56-9—it uses one SD port and one port on the analyzer, instead of using a full, two-port analyzer.

Figure 56-10 Fibre Channel Analyzer Using a Single SD Port

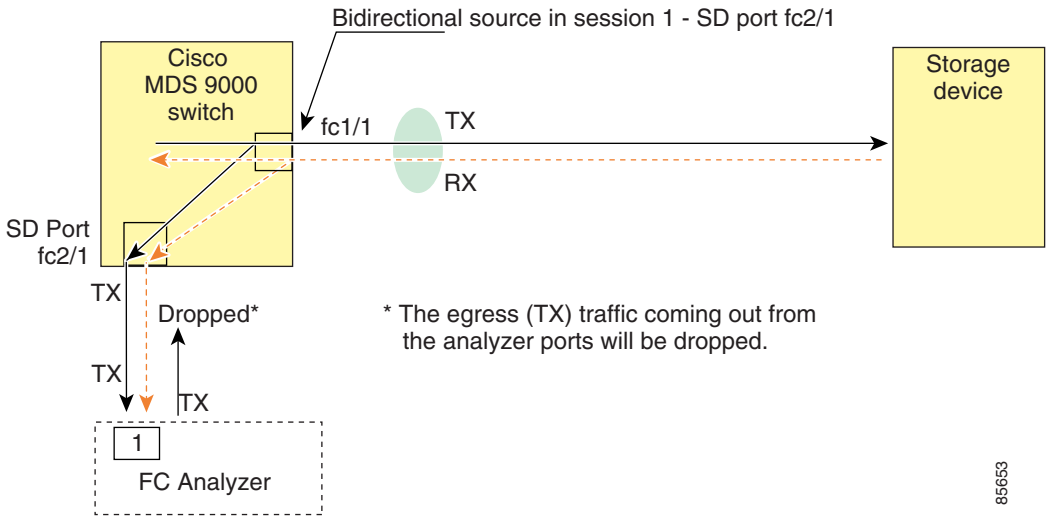

To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

Default Settings

Table 56-1 lists the default settings for SPAN parameters.

Table 56-1 Default SPAN Configuration Parameters

Parameters	Default
SPAN session	Active.
If filters are not specified	SPAN traffic includes traffic through a specific interface from all active VSANs.
Encapsulation	Disabled.
SD port	Output frame format is Fibre Channel.

CHAPTER 57

Configuring System Message Logging

This chapter describes configuration of system message logging on Cisco MDS 9000 Family switches. It includes the following sections:

About System Message Logging

The system message logging software saves messages in a log file or directs the messages to other devices. This feature provides you with the following capabilities:

- Provides logging information for monitoring and troubleshooting.
- Allows you to select the types of captured logging information.
- Allows you to select the destination server to forward the captured logging information.

By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. You can specify which system messages should be saved based on the type of facility (see Table 57-1) and the severity level (see Table 57-2). Messages are time-stamped to enhance real-time debugging and management.

The switch software saves system messages in a file that can be configured to save up to 4 MB. You can monitor system messages by clicking the **Events** tab on Fabric Manager or by choosing **Logs > Events > Current** on Device Manager. You can also monitor system messages remotely by accessing the switch through Telnet, SSH, or the console port, or by viewing the logs on a syslog server.

Note When the switch first initializes, the network is not connected until initialization completes. Therefore, messages are not redirected to a system message logging server for a few seconds.

Log messages are not saved across system reboots. However, a maximum of 100 log messages with a severity level of critical and below (levels 0, 1, and 2) are saved in NVRAM.

Table 57-1 describes some samples of the facilities supported by the system message logs.

Table 57-1 Internal Logging Facilities

Facility Keyword	Description	Standard or Cisco MDS Specific
acl	ACL manager	Cisco MDS 9000 Family specific
all	All facilities	Cisco MDS 9000 Family specific
auth	Authorization system	Standard
authpriv	Authorization (private) system	Standard
bootvar	Bootvar	Cisco MDS 9000 Family specific
callhome	Call Home	Cisco MDS 9000 Family specific
cron	Cron or at facility	Standard
daemon	System daemons	Standard
fcc	FCC	Cisco MDS 9000 Family specific
fcdomain	fcdomain	Cisco MDS 9000 Family specific
fcns	Name server	Cisco MDS 9000 Family specific
fcs	FCS	Cisco MDS 9000 Family specific
flogi	FLOGI	Cisco MDS 9000 Family specific
fspf	FSPF	Cisco MDS 9000 Family specific
ftp	File Transfer Protocol	Standard
ipconf	IP configuration	Cisco MDS 9000 Family specific
ipfc	IPFC	Cisco MDS 9000 Family specific
kernel	Kernel	Standard
local0 to local7	Locally defined messages	Standard
lpr	Line printer system	Standard
mail	Mail system	Standard
mcast	Multicast	Cisco MDS 9000 Family specific
module	Switching module	Cisco MDS 9000 Family specific
news	USENET news	Standard
ntp	NTP	Cisco MDS 9000 Family specific
platform	Platform manager	Cisco MDS 9000 Family specific
port	Port	Cisco MDS 9000 Family specific
port-channel	PortChannel	Cisco MDS 9000 Family specific
qos	QoS	Cisco MDS 9000 Family specific
rdl	RDL	Cisco MDS 9000 Family specific
rib	RIB	Cisco MDS 9000 Family specific
rscn	RSCN	Cisco MDS 9000 Family specific
securityd	Security	Cisco MDS 9000 Family specific
syslog	Internal system messages	Standard
sysmgr	System manager	Cisco MDS 9000 Family specific
tlport	TL port	Cisco MDS 9000 Family specific

Table 57-1 Internal Logging Facilities (continued)

Facility Keyword	Description	Standard or Cisco MDS Specific
user	User process	Standard
uucp	UNIX-to-UNIX Copy Program	Standard
vhbad	Virtual host base adapter daemon	Cisco MDS 9000 Family specific
vni	Virtual network interface	Cisco MDS 9000 Family specific
vrrp_cfg	VRRP configuration	Cisco MDS 9000 Family specific
vrrp_eng	VRRP engine	Cisco MDS 9000 Family specific
vsan	VSAN system messages	Cisco MDS 9000 Family specific
vshd	vshd	Cisco MDS 9000 Family specific
wwn	WWN manager	Cisco MDS 9000 Family specific
xbar	Xbar system messages	Cisco MDS 9000 Family specific
zone	Zone server	Cisco MDS 9000 Family specific

Table 57-2 describes the severity levels supported by the system message logs.

Table 57-2 Error Message Severity Levels

Level Keyword	Level	Description	System Message Definition
emergencies	0	System unusable	LOG_EMERG
alerts	1	Immediate action needed	LOG_ALERT
critical	2	Critical conditions	LOG_CRIT
errors	3	Error conditions	LOG_ERR
warnings	4	Warning conditions	LOG_WARNING
notifications	5	Normal but significant condition	LOG_NOTICE
informational	6	Informational messages only	LOG_INFO
debugging	7	Debugging messages	LOG_DEBUG

Note Refer to the *Cisco MDS 9000 Family System Messages Reference* for details on the error log message format.

System Message Logging Configuration

System logging messages are sent to the console based on the default (or configured) logging facility and severity values.

Message Logging Initiation

You can disable logging to the console or enable logging to a given Telnet or SSH session.

- When you disable or enable logging to a console session, that state is applied to all future console sessions. If you exit and log in again to a new session, the state is preserved.
- When you enable or disable logging to a Telnet or SSH session, that state is applied only to that session. If you exit and log in again to a new session, the state is not preserved.

To enable or disable the logging state for a Telnet or SSH session using Fabric Manager, follow these steps:

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Switches,** expand **Events and** select **SysLog** in the Physical Attributes pane.

You see the SysLog information in the Information pane.

Step 3 Click the **Switch Logging** tab.

You see the switch information shown in Figure 57-1.

Figure 57-1 Switch Logging Tab in Fabric Manager

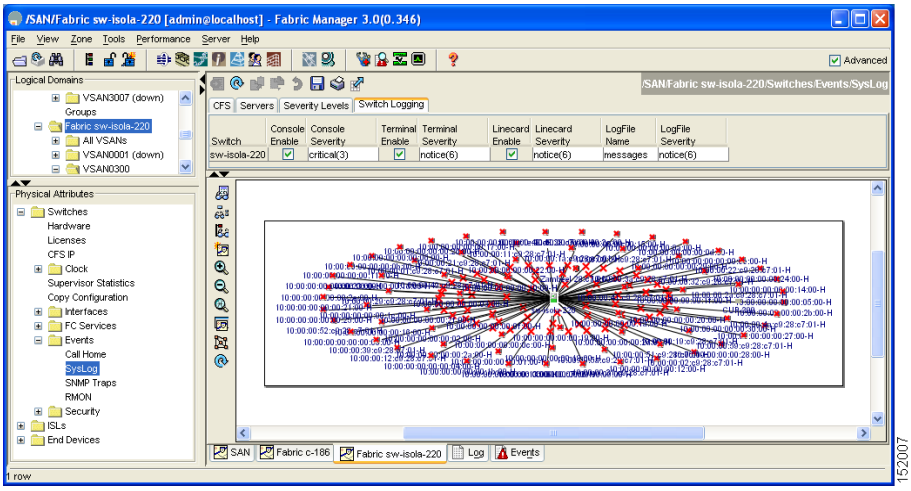

Step 4 Select a switch in the Information pane.

Step 5 Check (enable) or uncheck (disable) the **Console Enable** check box.

Step 6 Click the **Apply Changes** icon.

Console Severity Level

When logging is enabled for a console session (default), you can configure the severity levels of messages that appear on the console. The default severity for console logging is 2 (critical).

Tip The current critical (default) logging level is maintained if the console baud speed is 9600 baud (default). All attempts to change the console logging level generates an error message. To increase the logging level (above critical), you must change the console baud speed to 38400 baud.

To configure the severity level for a logging facility using Fabric Manager, follow these steps:

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Switches,** expand **Events and** select **SysLog** in the Physical Attributes pane.

You see the SysLog information in the Information pane.

Step 3 Click the **Switch Logging** tab.

You see the switch information shown in Figure 57-2.

Figure 57-2 Switch Logging Tab in Fabric Manager

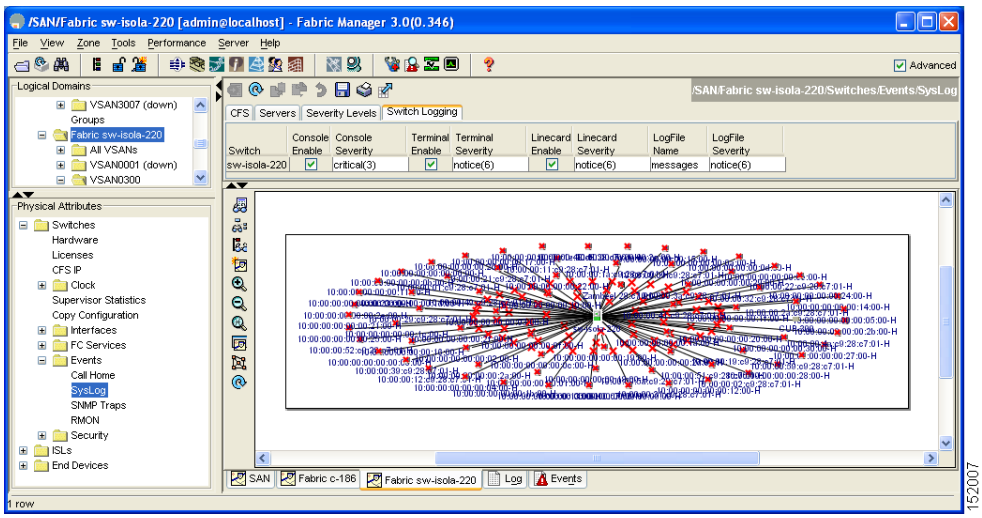

Step 4 Select a switch in the Information pane.

Step 5 Select a severity level from the Console Severity drop-down list in the row for that switch.

Step 6 Click the **Apply Changes** icon.

Module Logging

By default, logging is enabled at level 7 for all modules. You can enable or disable logging for each module at a specified level.

To configure the severity level for a logging facility, follow these steps:

Step 1 In Fabric Manager, expand **Switches,** expand **Events and** select **SysLog** in the Physical Attributes pane.

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

You see the switch information shown in Figure 57-3 or Figure 57-4.

Figure 57-3 Switch Logging Tab in Fabric Manager

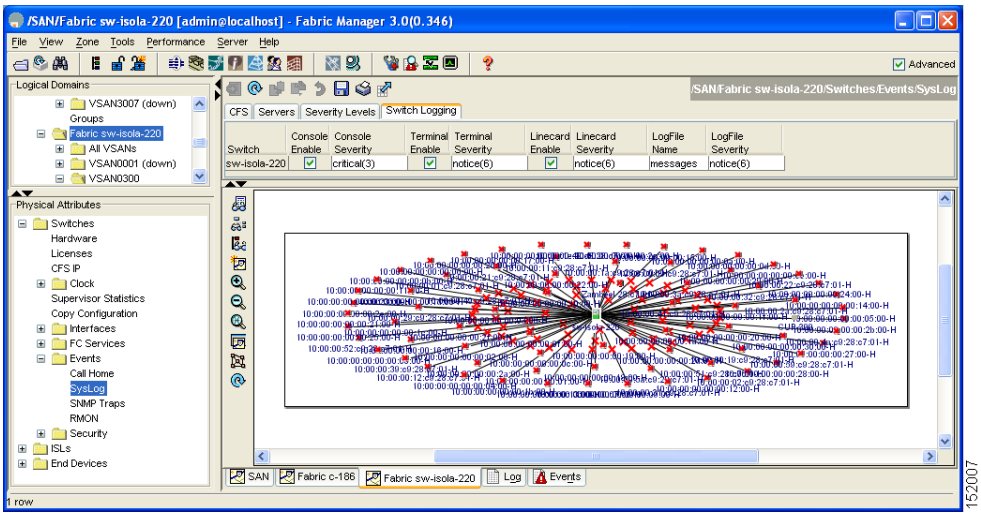

Figure 57-4 Switch Logging Tab in Device Manager

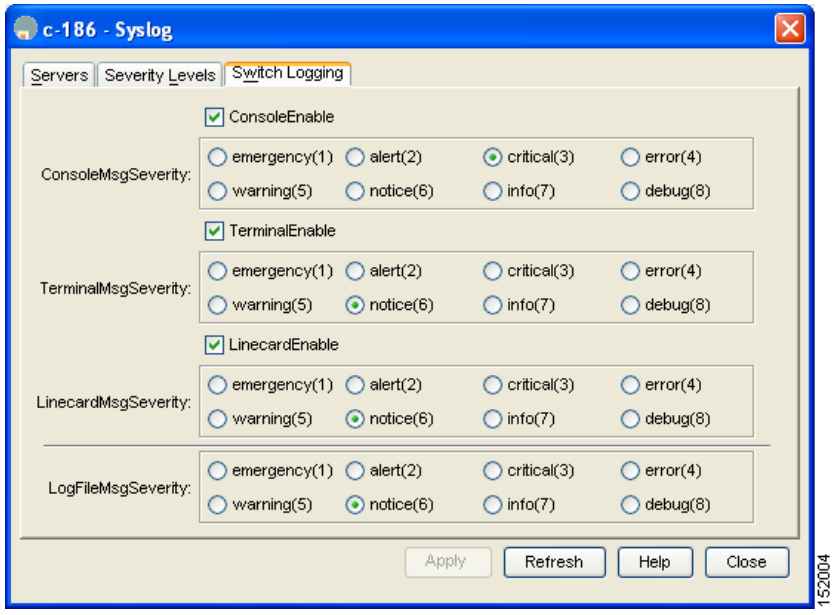

Step 2 Check the check boxes where you want message logging to occur (**ConsoleEnable**, **TerminalEnable**, **LineCardEnable**).

Step 3 Choose the message severity threshold from the **Console Severity** drop-down box for each switch in Fabric Manager (see Figure 57-3) or click the appropriate message severity level radio button in Device Manager (see Figure 57-4).

Step 4 Click the **Apply Changes** icon in Fabric Manager, or click **Apply** in Device Manager to save and apply your changes.

Log Files

Logging messages may be saved to a log file. You can configure the name of this file and restrict its size as required. The default log file name is messages. The file name can have up to 80 characters and the file size ranges from 4096 bytes to 4194304 bytes.

To send log messages to file using Fabric Manager, follow these steps:

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Switches,** expand **Events and** select **SysLog** in the Physical Attributes pane.

You see the SysLog information in the Information pane.

Step 3 Select a switch in the Information pane.

Step 4 Click the **Switch Logging** tab.

You see the information in Figure 57-5.

Figure 57-5 Switch Logging Tab in Fabric Manager

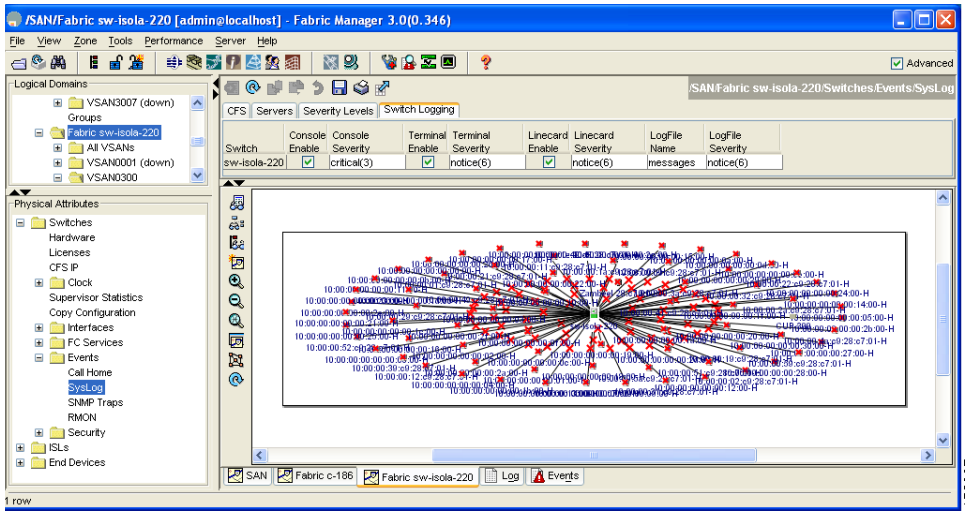

Step 5 Enter the name of the log file in the LogFile Name column in the row for that switch.

Step 6 Click the **Apply Changes** icon.

The configured log file is saved in the /var/log/external directory. The location of the log file cannot be changed.

System Message Logging Servers

You can configure a maximum of three system message logging servers.

To send log messages to a UNIX system message logging server, configure the system message logging daemon on a UNIX server. Log in as root, and follow these steps:

Step 1 Add the following line to the /etc/syslog.conf file.

```
local1.debug                    /var/log/myfile.log
```

Note Be sure to add five tab characters between **local1.debug** and **/var/log/*myfile*.log**. See entries in the /etc/syslog.conf file for further examples.

The switch sends messages according to the specified facility types and severity levels. The **local1** keyword specifies the UNIX logging facility used. The messages from the switch are generated by user processes. The **debug** keyword specifies the severity level of the condition being logged. You can set UNIX systems to receive all messages from the switch.

Step 2 Create the log file by entering these commands at the UNIX shell prompt:

```
$ touch /var/log/myfile.log
$ chmod 666 /var/log/myfile.log
```

Step 3 Make sure the system message logging daemon reads the new changes by entering this command:

```
$ kill -HUP ~cat /etc/syslog.pid~
```

Note Most tabs in the Information pane for features using CFS are dimmed until you click the CFS tab. The CFS tab shows which switches have CFS enabled and shows the master switch for this feature. Once the CFS tab is click, the other tabs in the Information pane that use CFS are activated.

You can configure a maximum of three syslog servers. One of these syslog servers should be Fabric Manager if you want to view system messages from the Event tab in Fabric Manager.

To configure system message logging servers, follow these steps:

Step 1 In Fabric Manager, Expand **Switches,** expand **Events** and select **SysLog** in the Physical Attributes pane, then click the **Servers** tab in the Information pane.

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

Figure 57-6 Servers Tab in Fabric Manager Syslog

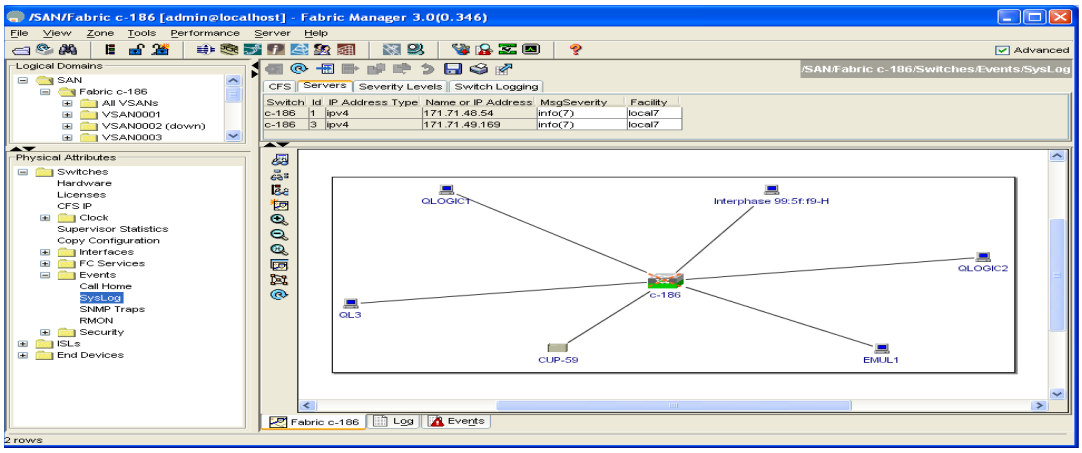

Figure 57-7 Servers Tab in Device Manager Syslog

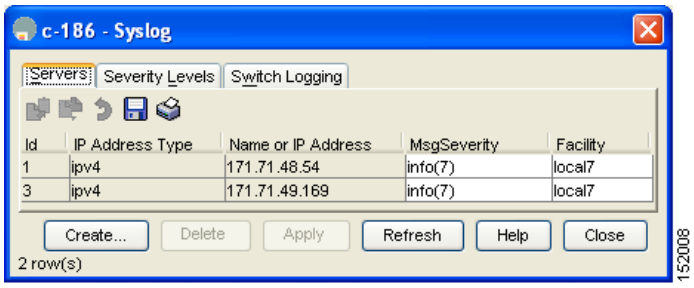

Step 2 Click the **Create Row icon** in Fabric Manager, or click **Create** in Device Manager (see Figure 57-7) to add a new syslog server.

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

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

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

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

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

- Hardware—Line card or power supply problems
- Link Incidents—FICON port condition changes
- Accounting—User change events

- Events—All other events

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

Verifying Syslog Servers from Fabric Manager Web Services

To verify the syslog servers remotely using Fabric Manager Web Services, follow these steps:

Step 1 Point your browser at the Fabric Manager Web Services server. See the "Launching Fabric Manager Web Services" section on page 6-5.

Step 2 Choose **Events > Syslog** to view the syslog server information for each switch. The columns in the table are sortable.

Outgoing System Message Logging Server Facilities

All system messages have a logging facility and a level. The logging facility can be thought of as *where* and the level can be thought of as *what*.

The single system message logging daemon (syslogd) sends the information based on the configured **facility** option. If no facility is specified, local7 is the default outgoing facility.

The internal facilities are listed in Table 57-1 and the outgoing logging facilities are listed in Table 57-3.

Table 57-3 Outgoing Logging Facilities

Facility Keyword	Description	Standard or Cisco MDS Specific
auth	Authorization system	Standard
authpriv	Authorization (private) system	Standard
cron	Cron or at facility	Standard
daemon	System daemons	Standard
ftp	File Transfer Protocol	Standard
kernel	Kernel	Standard
local0 to local7	Locally defined messages	Standard (local7 is the default)
lpr	Line printer system	Standard
mail	Mail system	Standard
news	USENET news	Standard
syslog	Internal system messages	Standard
user	User process	Standard
uucp	UNIX-to-UNIX Copy Program	Standard

Viewing Logs from Fabric Manager Web Services

To view system messages remotely using Fabric Manager Web Services, follow these steps:

Step 1 Point your browser at the Fabric Manager Web Services server. See the "Launching Fabric Manager Web Services" section on page 6-5.

Step 2 Click the **Events tab** followed by the **Details** to view the system messages. The columns in the events table are sortable. In addition, you can use the Filter button to limit the scope of messages within the table.

Viewing Logs from Device Manager

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

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

Default Settings

Table 57-4 lists the default settings for system message logging.

Table 57-4 Default System Message Log Settings

Parameters	Default
System message logging to the console	Enabled for messages at the critical severity level.
System message logging to Telnet sessions	Disabled.
Logging file size	4194304.
Log file name	Message (change to a name with up to 200 characters).
Logging server	Disabled.
Syslog server IP address	Not configured.
Number of servers	Three servers.
Server facility	Local 7.

CHAPTER 58

Configuring Call Home

Call Home provides e-mail-based notification of critical system events. A versatile range of message formats are available for optimal compatibility with pager services, standard e-mail, or XML-based automated parsing applications. Common uses of this feature may include direct paging of a network support engineer, e-mail notification to a Network Operations Center, and utilization of Cisco AutoNotify services for direct case generation with the Technical Assistance Center.

The Call Home feature provides message throttling capabilities. Periodic inventory messages, port syslog messages, and RMON alert messages are added to the list of deliverable Call Home messages. If required you can also use the Cisco Fabric Services application to distribute the Call Home configuration to all other switches in the fabric.

This chapter includes the following sections:

Call Home Features

The Call Home functionality is available directly through the Cisco MDS 9000 Family. It provides multiple Call Home profiles (also referred to as *Call Home destination profiles*), each with separate potential destinations. You can define your own destination profiles in addition to predefined profiles.

The Call Home function can even leverage support from Cisco Systems or another support partner. Flexible message delivery and format options make it easy to integrate specific support requirements.

The Call Home feature offers the following advantages:

- Fixed set of predefined alerts and trigger events on the switch.
- Automatic execution and attachment of relevant command output.
- Multiple message format options:
  - Short Text—Suitable for pagers or printed reports.
  - Plain Text—Full formatted message information suitable for human reading.
  - XML—Matching readable format using Extensible Markup Language (XML) and document type definitions (DTDs) named Messaging Markup Language (MML). The MML DTD is published on the Cisco.com website at http://www.cisco.com/. The XML format enables communication with the Cisco Systems Technical Assistance Center.
- Multiple concurrent message destinations. You can configure up to 50 e-mail destination addresses for each destination profile.
- Multiple message categories including system, environment, switching module hardware, supervisor module, hardware, inventory, syslog, RMON, and test.

Cisco AutoNotify

For those who have service contracts directly with Cisco Systems, automatic case generation with the Technical Assistance Center is possible by registering with the AutoNotify service. AutoNotify provides fast time to resolution of system problems by providing a direct notification path to Cisco customer support.

The AutoNotify feature requires several Call Home parameters to be configured, including certain contact information, e-mail server, and an XML destination profile as specified in the Service Activation document found on the Cisco.com web site at:
http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_3/service/serv332/ccmsrvs/sssrvact.htm

To configure a Cisco MDS 9000 Family switch to use the AutoNotify service, an XML destination profile must be configured to send messages to Cisco. Specific setup, activation, and e-mail address information is found on the Cisco.com web site at:
http://www.cisco.com/en/US/partner/products/hw/ps4159/ps4358/products_configuration_example09186a0080108e72.shtml

To register, the following items are required:

- The SMARTnet contract number covering your Cisco MDS 9000 Family switch.
- Your name, company address, your e-mail address, and your Cisco.com ID.

- The exact product number of your Cisco MDS 9000 Family switch. For example, some valid product numbers include: DS-C6509 and DS-C9216-K9.
- The serial number of your Cisco MDS 9000 Family switch. This can be obtained by looking at the serial number label on the back of the switch (next to the power supply).

The ContractID, CustomerID, SiteID, and SwitchPriority parameters are not required by the AutoNotify feature. They are only intended to be used as additional information by Cisco customers and service partners.

Call Home Configuration Process

The actual configuration of Call Home depends on how you intend to use the feature. Some points to consider include:

- An e-mail server and at least one destination profile (predefined or user-defined) must be configured. The destination profile(s) used depends on whether the receiving entity is a pager, e-mail, or automated service such as Cisco AutoNotify.
- The contact name (SNMP server contact), phone, and street address information must be configured before Call Home is enabled. This is required to determine the origin of messages received.
- The Cisco MDS 9000 switch must have IP connectivity to an e-mail server.
- If Cisco AutoNotify is used, an active service contract must cover the device being configured.

To configure Call Home, follow these steps:

Step 1 Assign contact information.

Step 2 Configure destination profiles.

Step 3 Associate one or more alert groups to each profile as required by your network.

Step 4 Enable or disable Call Home.

Step 5 Test Call Home messages.

Contact Information

It is mandatory for each switch to include e-mail, phone, and street address information. It is optional to include the contract ID, customer ID, site ID, and switch priority information.

To assign the contact information, follow these steps:

Step 1 In the Fabric Manager Physical Attributes pane, expand **Switches,** expand **Events** and select **Call Home**. You see the Call Home in the Information pane. See Figure 58-1.

Figure 58-1 Call Home in Fabric Manager

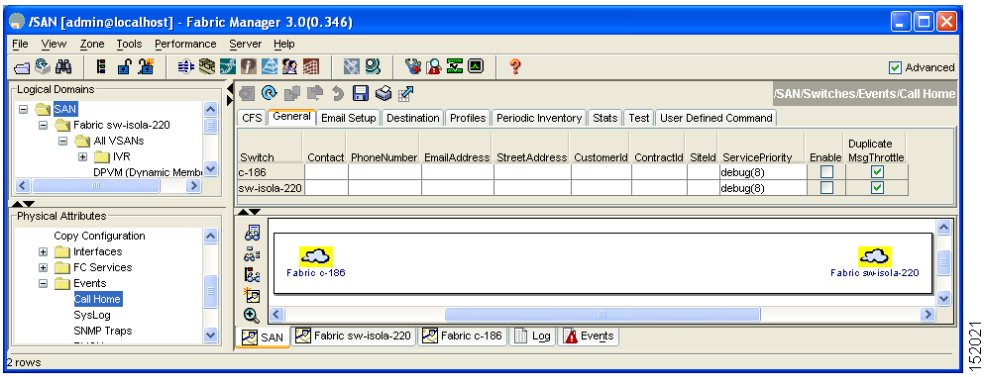

In Device Manager, click **Admin > Events > Call Home**. See Figure 58-2.

Figure 58-2 Call Home in Device Manager

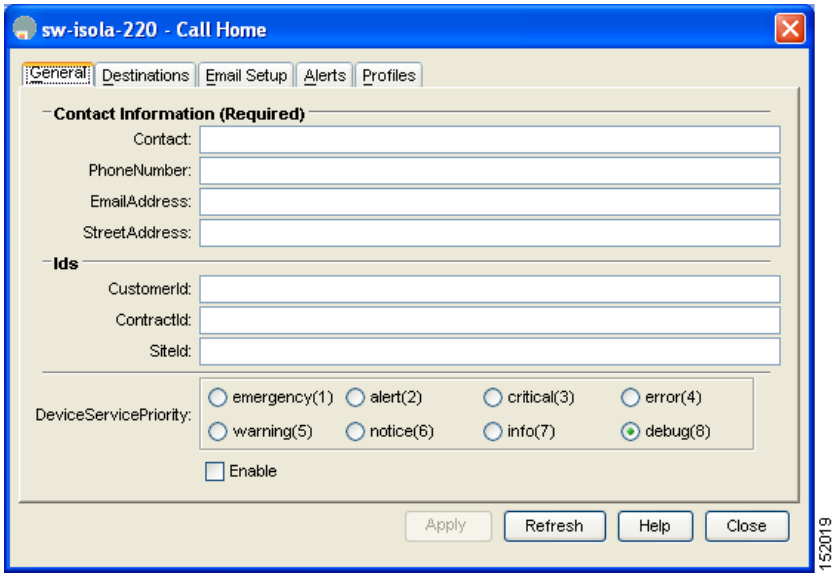

Step 2 Click the **General** tab, then assign contact information and enable the Call Home feature. Call Home is not enabled by default. You must enter an e-mail address that identifies the source of Call Home notifications.

Step 3 Click the **Destination(s)** tab to configure the destination e-mail addresses for Call Home notifications. You can identify one or more e-mail addresses that will receive Call Home notifications.

Step 4 Click the **Email Setup** tab to identify the SMTP server. Identify a message server to which your switch has access. This message server will forward the Call Home notifications to the destinations.

Step 5 In Fabric Manager, click the **Apply Changes** icon. In Device Manager, click **Apply**.

Destination Profiles

A destination profile contains the required delivery information for an alert notification. Destination profiles are typically configured by the network administrator. At least one destination profile is required. You can configure multiple destination profiles of one or more types.

You can use one of the predefined destination profiles or define a desired profile. If you define a new profile, you must assign a profile name.

Note If you use the Cisco AutoNotify service, the XML destination profile is required (see http://www.cisco.com/warp/customer/cc/serv/mkt/sup/tsssv/opmsup/smton/anoti_ds.htm).

- Profile name—A string that uniquely identifies each user-defined destination profile and is limited to 32 alphanumeric characters. The format options for a user-defined destination profile are full-txt, short-txt, or XML (default).
- Destination address—The actual address, pertinent to the transport mechanism, to which the alert should be sent.
- Message formatting—The message format used for sending the alert (full text, short text, or XML).

To configure predefined destination profile messaging options using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **Events**, and select **Call Home** in the Physical Attributes pane then click the **Profiles** tab in the Information pane.

You see the Call Home profiles for multiple switches shown in Figure 58-3.

Figure 58-3 Call Home Profiles for Multiple Switches

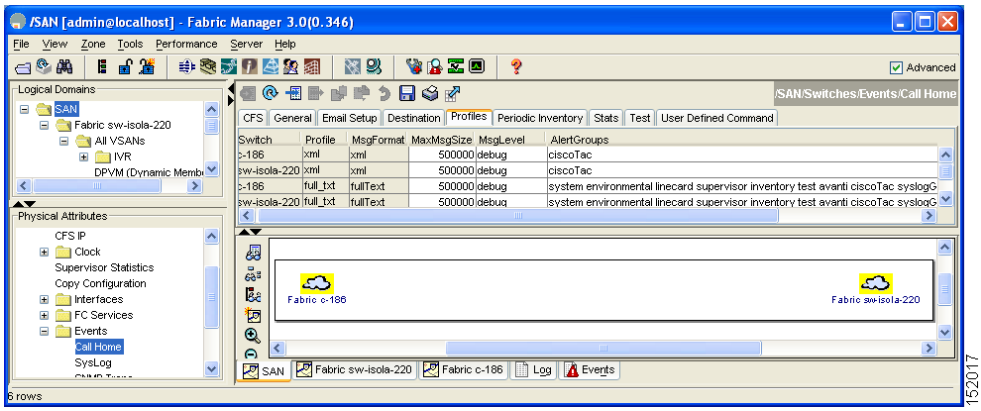

Step 2 Set the profile name, message format, message size, and severity level.

Step 3 Click in the Alert Groups column and select or remove an alert group.

Step 4 Click the **Apply Changes** icon to create this profile on the selected switches.

To configure a new destination-profile (and related parameters), follow these steps:

Step 1 Expand **Switches**, expand **Events**, and select **Call Home** in the Physical Attributes pane then click the **Profiles** tab in the Information pane.

You see Call Home profiles for multiple switches.

Figure 58-4 Call Home Profiles for Multiple Switches

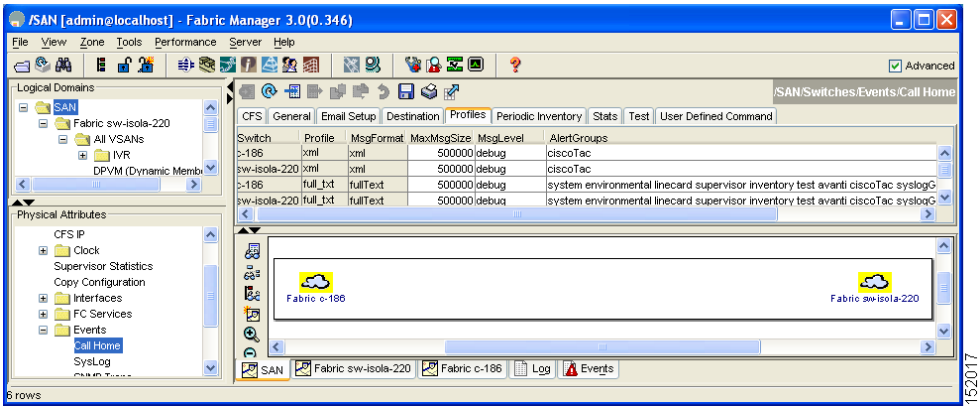

Step 2 Click the **Create Row** icon to add a new profile.

Step 3 Set the profile name, message format, size, and severity level.

Step 4 Click an alert group and select each group from the drop-down list that you want sent in this profile.

Step 5 Click **Create** to create this profile on the selected switches.

Alert Groups

An alert group is a predefined subset of Call Home alerts supported in all switches in the Cisco MDS 9000 Family. Different types of Call Home alerts are grouped into different alert groups depending on their type. You can associate one or more alert groups to each profile as required by your network.

The alert group allows you to select the set of Call Home alerts to be received by a destination profile (predefined or user-defined). You can associate multiple alert groups with a destination profile.

Note A Call Home alert is sent to e-mail destinations in a destination profile only if that Call Home alert belongs to one of the alert groups associated with that destination profile.

To associate an alert group with a destination profile, follow these steps:

Step 1 Expand **Switches**, expand **Events**, and select **Call Home** in the Physical Attributes pane then click the **Profiles** tab in the Information pane.

You see the Call Home profiles for multiple switches shown in Figure 58-5.

Figure 58-5 Call Home Profiles for Multiple Switches

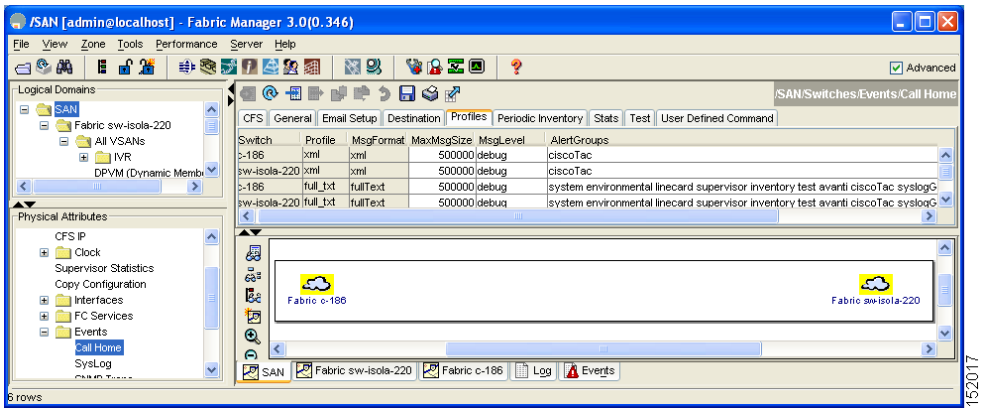

Step 2 Click the **Alert Groups** column in the row for the profile you want to associate.

You see the alert groups drop-down menu shown in Figure 58-6.

Figure 58-6 Alert Groups Drop-down Menu

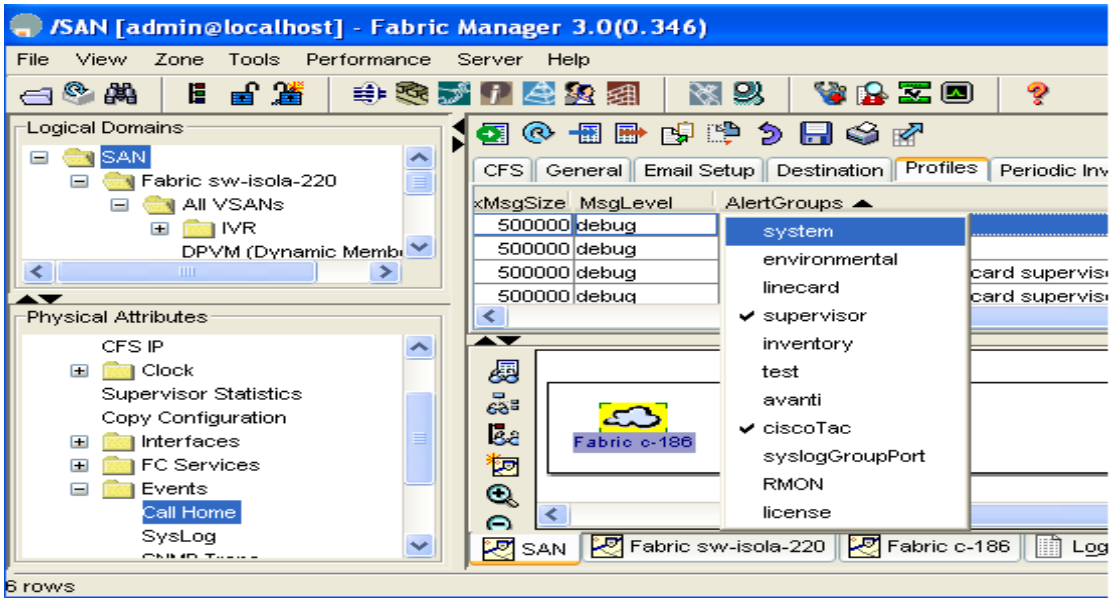

Step 3 Click an alert group to select it for association.

You see a check next to that alert group. To deselect it and remove the check, click it again.

Step 4 Click the **Apply Changes** icon.

Customized Alert Group Messages

The predefined Call Home alert groups generate notification messages when certain events occur on the switch. You can customize predefined alert groups to execute additional show commands when specific events occur. The output from these additional show commands is included in the notification message along with that of the predefined show commands.

Note You can assign a maximum of five show commands to an alert group. Only show commands can be assigned to an alert group.

Note Customized show commands are only supported for full text and XML alert groups. Short text alert groups (short-txt-destination) do not support customized show commands because they only allow 128 bytes of text.

To assign **show** commands to be executed when an alert is sent, you must associate the commands with the alert group. When an alert is sent, Call Home associates the alert group with an alert type and attaches the output of the **show** commands to the alert message.

Note Make sure the destination profiles for the non-Cisco-TAC alert group, with a predefined **show** command, and the Cisco-TAC alert group are not the same.

To customize Call Home alert group messages using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **Events**, and select **Call Home** in the Physical Attributes pane, then click the **User Defined Command** tab in the Information pane.

Step 2 Click the **Create Row** icon.

You see the User Defined Command information shown in Figure 58-7.

Figure 58-7 User Defined Command Dialog Box

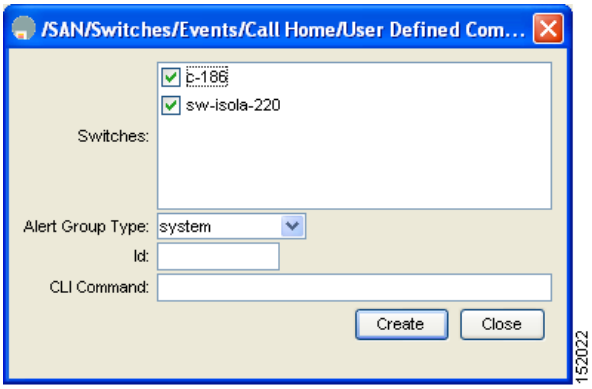

Step 3 Check the check boxes in front of the switches from which you want to receive alerts.

Step 4 Select the alert group type from the Alert Group Type drop-down list.

Step 5 Select the ID (1-5) of the CLI command. The ID is used to keep track of the messages.

Step 6 Enter the CLI **show** command in the **CLI Command** field.

Step 7 Click **Create**.

Step 8 Repeat Steps 3-7 for each command you want to associate with the profile.

Step 9 Click **Close** to close the dialog box.

Call Home Message Levels

The Call Home message level feature allows you to filter messages based on their level of urgency. Each destination profile (predefined and user-defined) is associated with a Call Home message level threshold. Any message with a value lower that the urgency threshold is not sent. The urgency level ranges from 0 (lowest level of urgency) to 9 (highest level of urgency), and the default is 0 (all messages are sent).

Note Call Home severity levels are not the same as system message logging severity levels.

To set the message level for each profile for Call Home, follow these steps:

Step 1 In Fabric Manager, expand the **Switches** folder in the Physical Attributes pane, expand **Events** then select **Call Home**.

You see the Call Home dialog box in the Information pane.

In Device Manager, choose **Admin > Events > Call Home**.

Step 2 Click the **Profiles** tab.

You see the Call Home profiles shown in Figure 58-8.

Figure 58-8 Call Home Profiles

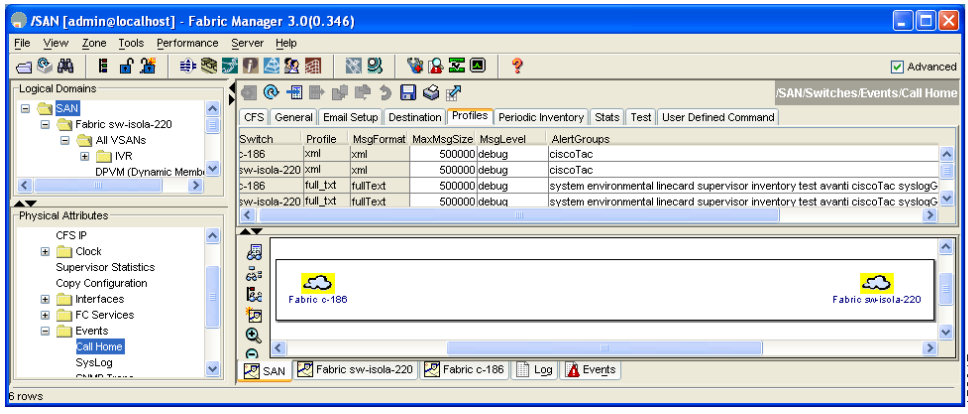

Step 3 Set a message level for each switch using the drop-down menu in the MsgLevel column.

Step 4 Click the **Apply Changes** icon to save your changes or click the **Undo Changes** icon to cancel your changes.

Syslog-Based Alerts

You can configure the switch to send certain syslog messages as Call Home messages. The syslog-group-port alert group selects syslog messages for the port facility. The Call Home application maps the syslog severity level to the corresponding Call Home severity level. For example, if you select level 5 for the Call Home message level, syslog messages at levels 0, 1, and 2 are included in the Call Home log.

Note Call Home does not change the syslog message level in the message text. The syslog message texts in the Call Home log appear as they are described in the *Cisco MDS 9000 Family System Messages Guide*.

Whenever a syslog message is generated, the Call Home application sends a Call Home Message depending on the mapping between the destination profile and the alert group mapping and based on the severity level of the generated syslog message. To receive a syslog-based Call Home alert, you must associate a destination profile with the syslog alert groups (currently there is only one syslog alert group—syslog-group-port) and configure the appropriate message level.

To configure the syslog-group-port using Fabric Manager, follow these steps:

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Switches**, expand **Events**, and select **Call Home** in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Select the **Profiles** tab.

You see the Call Home profiles shown in Figure 58-9.

Figure 58-9 Call Home Profiles

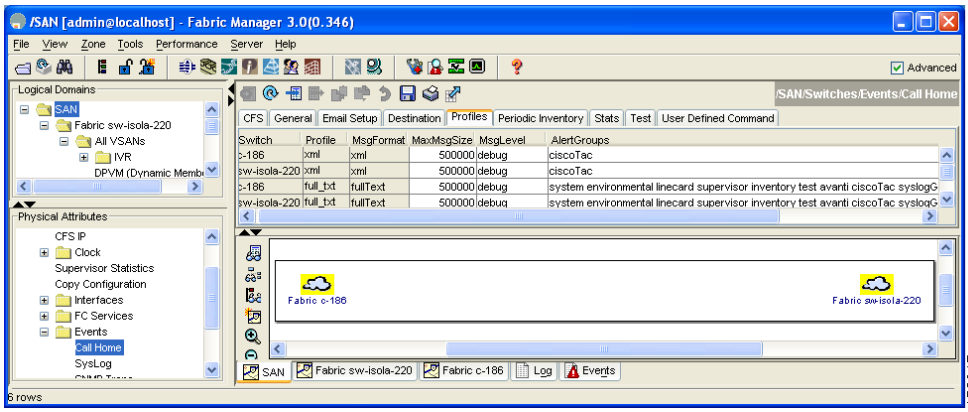

Step 4 Click the **Create Row** icon.

You see the Create Call Home Profile dialog box.

Step 5 Select the switches for which you want to send alerts.

Step 6 Enter the name of the profile in the Name field.

Step 7 Choose the message format, message size, and message severity level.

Step 8 Check the **syslogGroupPort** check box in the AlertGroups section.

Step 9 Click **Create** to create the profile for the syslog-based alerts.

Step 10 Close the dialog box.

RMON-Based Alerts

You can configure the switch to send Call Home notifications corresponding to RMON alert triggers. All RMON-based Call Home messages have their message level set to NOTIFY (2). The RMON alert group is defined for all RMON-based Call Home alerts. To receive an RMON-based Call Home alert, you must associate a destination profile with the RMON alert group.

To configure RMON alert groups using Fabric Manager, follow these steps:

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Switches**, expand **Events**, and select **Call Home** in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Select the **Profiles** tab.

You see the Call Home profiles shown in Figure 58-10.

Figure 58-10 Call Home Profiles

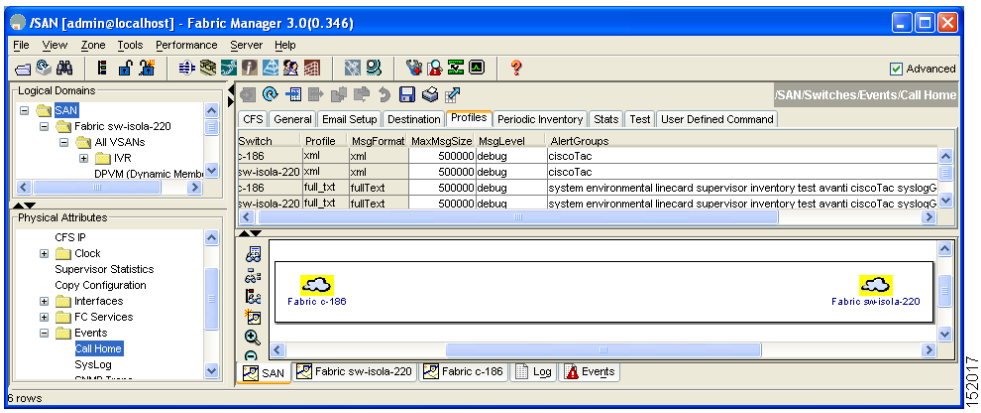

Step 4 Select the **Create Row** icon.

You see the Create Call Home Profile dialog box.

Step 5 Select switches for which to send alerts.

Step 6 Enter the name of the profile.

Step 7 Select the message format, message size, and message severity level.

Step 8 Check the **RMON** check box in the AlertGroups section.

Step 9 Click **Create** to create the profile for the RMON-based alerts.

Step 10 Close the dialog box.

E-Mail Options

You can configure the from, reply-to, and return-receipt e-mail addresses. While most e-mail address configurations are optional, you must configure the SMTP server address for the Call Home functionality to work.

Configuring General E-Mail Options

To configure general e-mail options and the SMTP server and port using Fabric Manager, follow these steps:

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Switches**, expand **Events**, and select **Call Home** in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Click the **Email Setup** tab.

Figure 58-11 Call Home Email Setup

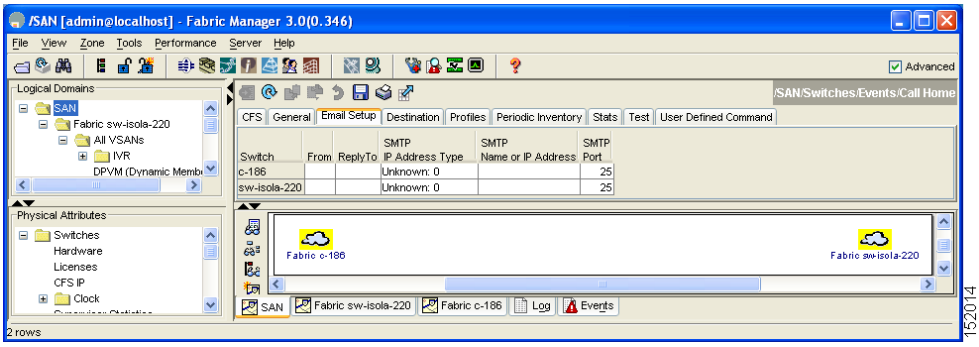

Step 4 Select a switch in the Information pane.

Step 5 Enter the general e-mail information.

Step 6 Enter the SMTP server IP address type, IP address or name, and port.

Step 7 Click the **Apply Changes** icon to update the e-mail options.

Periodic Inventory Notification

You can configure the switch to periodically send a message with an inventory of all the software services currently enabled and running on the switch along with hardware inventory information. The inventory is modified each time the switch is restarted nondisruptively.

By default, this feature is disabled in all switches in the Cisco MDS 9000 Family. When you enable this feature without configuring an interval value, the Call Home message is sent every 7 days. This value ranges from 1 to 30 days.

To enable periodic inventory notification in a Cisco MDS 9000 Family switch using Fabric Manager, follow these steps:

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Switches**, expand **Events**, and select **Call Home** in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Click the **Periodic Inventory** tab.

You see the Call Home periodic inventory information shown in Figure 58-12.

Figure 58-12 Call Home Periodic Inventory

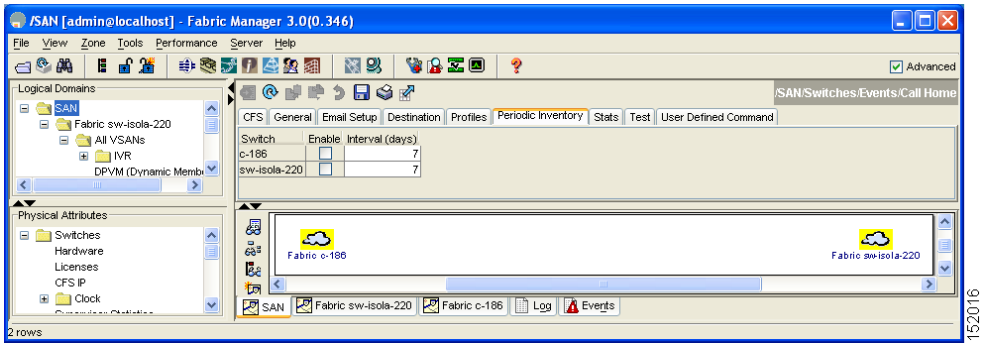

Step 4 Select a switch in the Information pane.

Step 5 Check the **Enable** check box.

Step 6 Enter the number of days for which you want the inventory checked.

Step 7 Click the **Apply Changes** icon.

Duplicate Message Throttle

You can configure a throttling mechanism to limit the number of Call Home messages received for the same event. If the same message is sent multiple times from the switch within a short period of time, you may be swamped with a large number of duplicate messages.

By default, this feature is enabled in all switches in the Cisco MDS 9000 Family. When enabled, if the number of messages sent exceeds the maximum limit of 30 messages within the 2-hour time frame, then further messages for that alert type are discarded within that time frame. You cannot modify the time frame or the message counter limit.

If 2 hours have elapsed since the first such message was sent and a new message has to be sent, then the new message is sent and the time frame is reset to the time when the new message was sent and the count is reset to 1.

To enable message throttling in a Cisco MDS 9000 Family switch using Fabric Manager, follow these steps:

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Switches**, expand **Events**, and select **Call Home** in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Select the **General** tab.

You see the information shown in Figure 58-13.

Figure 58-13 Call Home General Information

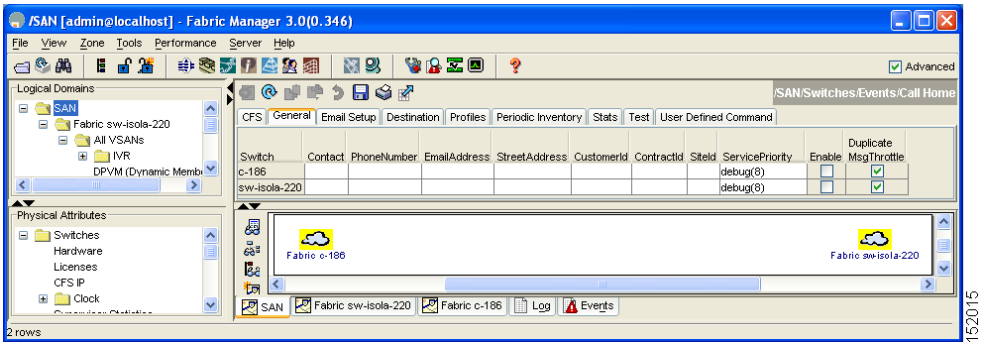

Step 4 Select a switch in the Information pane.

Step 5 Check the **Duplicate Message Throttle** check box.

Step 6 Click the **Apply Changes** icon.

Call Home Enable Function

Once you have configured the contact information, you must enable the Call Home function.

To enable the Call Home function using Fabric Manager, follow these steps:

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Switches**, expand **Events**, and select **Call Home** in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Select the **General** tab.

You see the information shown in Figure 58-14.

Figure 58-14 Call Home General Information

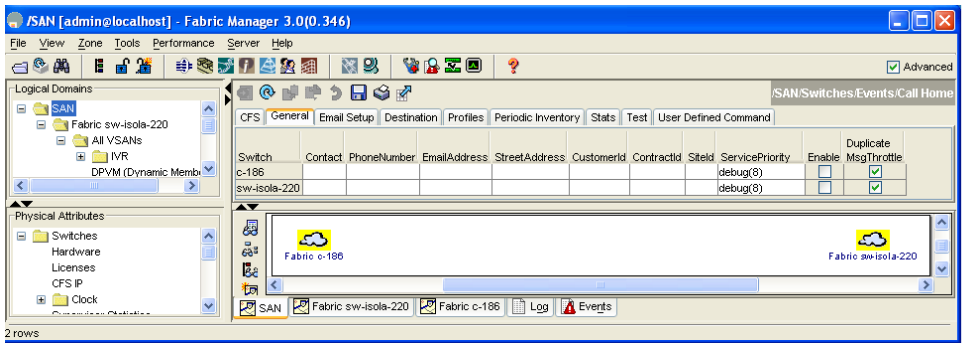

Step 4 Select a switch in the Information pane.

Step 5 Check the **Enable** check box.

Step 6 Click the **Apply Changes** icon.

Call Home Configuration Distribution

You can enable fabric distribution for all Cisco MDS switches in the fabric. When you perform Call Home configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.

Note The Switch priority and the Syscontact name are not distributed.

You automatically acquire a fabric-wide lock when you issue the first configuration operation after you enabled distribution in a switch. The Call Home application uses the effective and pending database model to store or commit the configuration changes. When you commit the configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. After making the configuration changes, you can choose to discard the changes by aborting the changes instead of committing them. In either case, the lock is released. See Chapter 12, “Using the CFS Infrastructure” for more information on the CFS application.

To enable Call Home fabric distribution, follow these steps:

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Switches**, expand **Events**, and select **Call Home** in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Select the **CFS** tab.

You see the CFS information for Call Home shown in Figure 58-15.

Figure 58-15 Call Home CFS Information

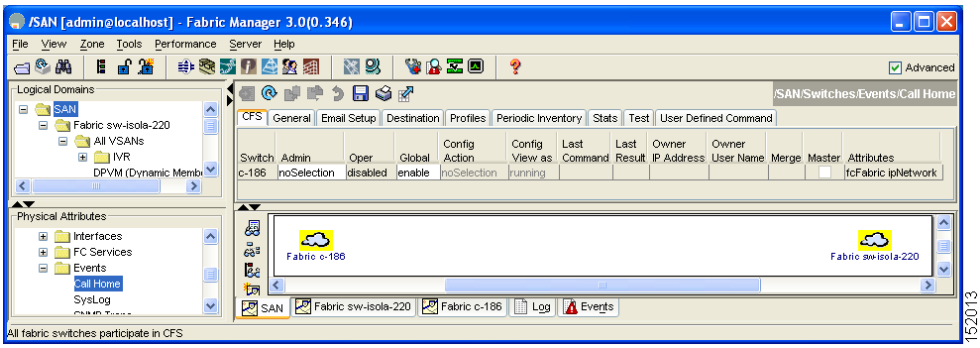

Step 4 Select a switch in the Information pane.

Step 5 Select **Enable** from the drop-down list in the Admin column in the row for that switch.

Step 6 Click the **Apply Changes** icon to commit the changes.

Fabric Lock Override

If you have performed a Call Home task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.

Tip The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

Database Merge Guidelines

See the "CFS Merge Support" section on page 12-9 for detailed concepts.

When merging two Call Home databases, follow these guidelines:

- Be aware that the merged database contains the following information:
  - A superset of all the destination profiles from the dominant and subordinate switches take part in the merge protocol.
  - The e-mail addresses and alert groups for the destination profiles.
  - Other configuration information (for example, message throttling, periodic inventory) from the switch that existed in the dominant switch before the merge.
- Verify that two destination profiles do not have the same name (even if they have different configuration information) on the subordinate and dominant switches. If they do contain the same name, the merge operation will fail. You must then modify or delete the conflicting destination profile on the required switch.

Call Home Communications Test

To test the Call Home function and simulate a message generation using Fabric Manager, follow these steps:

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Switches**, expand **Events**, and select **Call Home** in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Click the **Test** tab.

You see the configured tests for the switch and the status of the last testing.

Step 4 Select a switch in the Information pane.

Step 5 Select **test** or **testWithInventory** from the TestAction drop-down list in the row for that switch.

Step 6 Click the **Apply Changes** icon to run the test.

Configuring EMC Call Home

This feature is configured using Fabric Manager Web Services or by editing the server.properties file. The documentation for configuring EMC Call Home using Fabric Manager Web Services is contained in the Web Services **Admin > Configure > Preferences** web page. The documentation for configuring EMC Call Home by editing the server.properties file is contained within the server.properties file.

Sample Syslog Alert Notification in Full-txt Format

```
source:MDS9000
Switch Priority:7
Device Id:DS-C9506@C@FG@07120011
Customer Id:basu
Contract Id:123
Site Id:Bangalore
Server Id:DS-C9506@C@FG@07120011
Time of Event:2004-10-08T11:10:44
Message Name:SYSLOG_ALERT
Message Type:Syslog
Severity Level:2
System Name:10.76.100.177
Contact Name:Basavaraj B
Contact Email:bbendige@cisco.com
Contact Phone:+91-80-310-1718
Street Address:#71 , Miller's Road
Event Description:2004 Oct  8 11:10:44 10.76.100.177 %PORT-5-IF_TRUNK_UP: %$VSAN 1%$
Interface fc2/5, vsan 1 is up

syslog_facility:PORT
start chassis information:
Affected Chassis:DS-C9506
Affected Chassis Serial Number:FG@07120011
Affected Chassis Hardware Version:0.104
Affected Chassis Software Version:2.0(1)
Affected Chassis Part No:73-8607-01
```

```
end chassis information:
```

Sample Syslog Alert Notification in XML Format

```
X-Mozilla-Status2: 02000000
Return-Path: <tester@cisco.com>
...

<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<!DOCTYPE mml SYSTEM "mml10.dtd">
<!--
Alert:SYSLOG_ALERT
-->
<mml>
<header>
<time>2004-09-30T06:12:36</time>
<name>SYSLOG_ALERT</name>
<type>Syslog</type>
<level>2</level>
<source>MDS9000</source>
<priority>7</priority>
<deviceId>DS-C9506@C@FOX0712S00H</deviceId>
<custId>911</custId>
<contractId>33445</contractId>
<siteId>91111</siteId>
<serverId>DS-C9506@C@FOX0712S00H</serverId>
</header>
<body>
<msgDesc>2004 Sep 30 06:12:36 switch186 %PORT-5-IF_UP: %$VSAN 2000%$ Interface fc1/10 is
up in mode FL
</msgDesc>
<sysName>switch186</sysName>
<sysContact>USA</sysContact>
<sysContactEmail>billgates@microsoft.com</sysContactEmail>
<sysContactPhoneNumber>+91-080-8888888</sysContactPhoneNumber>
<sysStreetAddress>91</sysStreetAddress>
<chassis>
<name>DS-C9506</name>
<serialNo>FOX0712S00H</serialNo>
<partNo>73-8697-01</partNo>
<hwVersion>0.104</hwVersion>
<swVersion>2.0(1)</swVersion>
</chassis>
<nvp>
<name>syslog_facility</name>
<value>PORT</value>
</nvp>
</body>
</mml>
```

Sample RMON Notification in XML Format

```
Return-Path: <tester@cisco.com>
...
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<!DOCTYPE mml SYSTEM "mml10.dtd">
<!--
Alert:RMON_ALERT
-->
```

```
<mml>
<header>
<time>2004-10-12T04:59:13</time>
<name>RMON_ALERT</name>
<type>RMON</type>
<level>2</level>
<source>MDS9000</source>
<priority>3</priority>
<deviceId>DS-C9506@C@FOX0712S00H</deviceId>
<custId>0</custId>
<contractId>u</contractId>
<siteId>&amp;amp;</siteId>
<serverId>DS-C9506@C@FOX0712S00H</serverId>
</header>
<body>
<msgDesc>rlaxmina-w2k07</msgDesc>
<sysName>switch186</sysName>
<sysContact>USA</sysContact>
<sysContactEmail>billgates@microsoft.com</sysContactEmail>
<sysContactPhoneNumber>+91-080-000000</sysContactPhoneNumber>
<sysStreetAddress>91</sysStreetAddress>
<chassis>
<name>DS-C9506</name>
<serialNo>FOX0712S00H</serialNo>
<partNo>73-8697-01</partNo>
<hwVersion>0.104</hwVersion>
<swVersion>2.0(1)</swVersion>
</chassis>
<nvp>
<name>ThresholdType</name>
<value>RisingThreshold</value>
</nvp>
<nvp>
<name>ThresholdValue</name>
<value>0</value>
</nvp>
<nvp>
<name>AlarmValue</name>
<value>0</value>
</nvp>
</body>
</mml>
```

Default Settings

Table 58-1 lists the default Call Home settings.

Table 58-1 Default Call Home Settings

Parameters	Default
Destination message size for a message sent in full text format.	500,000.
Destination message size for a message sent in XML format.	500,000.
Destination message size for a message sent in short text format.	4,000.
DNS or IP address of the SMTP server to reach the server if no port is specified.	25.
Alert group association with profile.	All.

Table 58-1 Default Call Home Settings (continued)

Parameters	Default
Format type.	XML.
Call Home message level.	0 (zero).

Event Triggers

This section discusses Call Home trigger events. Trigger events are divided into categories, with each category assigned CLI commands to execute when the event occurs. The command output is included in the transmitted message. Table 58-2 lists the trigger events.

Table 58-2 Event Triggers

Event	Alert Group	Event Name	Description	Call Home Message Level
Call Home	System and CISCO_TAC	SW_CRASH	A software process has crashed with a stateless restart, indicating an interruption of a service.	5
	System and CISCO_TAC	SW_SYSTEM_INCONSISTENT	Inconsistency detected in software or file system.	5
	Environmental and CISCO_TAC	TEMPERATURE_ALARM	Thermal sensor indicates temperature reached operating threshold.	6
		POWER_SUPPLY_FAILURE	Power supply failed.	6
		FAN_FAILURE	Cooling fan has failed.	5
	Switching module and CISCO_TAC	LINECARD_FAILURE	Switching module operation failed.	7
		POWER_UP_DIAGNOSTICS_FAILURE	Switching module failed power-up diagnostics.	7
	Line Card Hardware and CISCO_TAC	PORT_FAILURE	Hardware failure of interface port(s).	6
	Line Card Hardware, Supervisor Hardware, and CISCO_TAC	BOOTFLASH_FAILURE	Failure of boot compact Flash card.	6
	Supervisor module and CISCO_TAC	SUP_FAILURE	Supervisor module operation failed.	7
		POWER_UP_DIAGNOSTICS_FAILURE	Supervisor module failed power-up diagnostics.	7
	Supervisor Hardware and CISCO_TAC	INBAND_FAILURE	Failure of in-band communications path.	7
	Supervisor Hardware and CISCO_TAC	EOBC_FAILURE	Ethernet out-of-band channel communications failure.	6
	Supervisor Hardware and CISCO_TAC	MGMT_PORT_FAILURE	Hardware failure of management Ethernet port.	5
	License	LICENSE_VIOLATION	Feature in use is not licensed, and are turned off after grace period expiration.	6

Table 58-2 Event Triggers (continued)

Event	Alert Group	Event Name	Description	Call Home Message Level
Inventory	Inventory and CISCO_TAC	COLD_BOOT	Switch is powered up and reset to a cold boot sequence.	2
		HARDWARE_INSERTION	New piece of hardware inserted into the chassis.	2
		HARDWARE_REMOVAL	Hardware removed from the chassis.	2
Test	Test and CISCO_TAC	TEST	User generated test.	2
Port syslog	Syslog-group-port	SYSLOG_ALERT	Syslog messages corresponding to the port facility.	2
RMON	RMON	RMON_ALERT	RMON alert trigger messages.	2

Table 58-3 lists event categories and command outputs.

Table 58-3 Event Categories and Executed Commands

Event Category	Description	Executed Commands
System	Events generated by failure of a software system that is critical to unit operation.	**show tech-support** **show system redundancy status**
Environmental	Events related to power, fan, and environment sensing elements such as temperature alarms.	**show module** **show environment**
Switching module hardware	Events related to standard or intelligent switching modules.	**show tech-support**
Supervisor hardware	Events related to supervisor modules.	**show tech-support**
Inventory	Inventory status is provided whenever a unit is cold booted, or when FRUs are inserted or removed. This is considered a noncritical event, and the information is used for status and entitlement.	**show version**
Test	User generated test message.	**show version**

Call Home Message Levels

Call Home messages (sent for syslog alert groups) have the syslog severity level mapped to the Call Home message level.

This section discusses the severity levels for a Call Home message when using one or more switches in the Cisco MDS 9000 Family. Call Home message levels are preassigned per event type.

Severity levels range from 0 to 9, with 9 having the highest urgency. Each syslog level has keywords and a corresponding syslog level as listed in Table 58-4.

Note Call Home does not change the syslog message level in the message text. The syslog message texts in the Call Home log appear as they are described in the *Cisco MDS 9000 Family System Messages Guide.*

Note Call Home severity levels are not the same as system message logging severity levels (see Chapter 57, "Configuring System Message Logging").

Table 58-4 Severity and Syslog Level Mapping

Call Home Level	Keyword Used	Syslog Level	Description
Catastrophic (9)	**Catastrophic**	N/A	Network wide catastrophic failure.
Disaster (8)	**Disaster**	N/A	Significant network impact.
Fatal (7)	**Fatal**	Emergency (0)	System is unusable.
Critical (6)	**Critical**	Alert (1)	Critical conditions, immediate attention needed.
Major (5)	**Major**	Critical (2)	Major conditions.
Minor (4)	**Minor**	Error (3)	Minor conditions.
Warning (3)	**Warning**	Warning (4)	Warning conditions.
Notify (2)	**Notification**	Notice (5)	Basic notification and informational messages. Possibly independently insignificant.
Normal (1)	**Normal**	Information (6)	Normal event signifying return to normal state.
Debug (0)	**Debugging**	Debug (7)	Debugging messages.

Message Contents

The following contact information can be configured on the switch:

- Name of the contact person
- Phone number of the contact person
- E-mail address of the contact person
- Mailing address to which replacement parts must be shipped, if required
- Site ID of the network where the site is deployed
- Contract ID to identify the service contract of the customer with the service provider

Table 58-5 describes the short text formatting option for all message typesFigure 58-15.

Table 58-5 Short Text Messages

Data Item	Description
Device identification	Configured device name
Date/time stamp	Time stamp of the triggering event
Error isolation message	Plain English description of triggering event
Alarm urgency level	Error level such as that applied to system message

Table 58-6, Table 58-7, and Table 58-8 display the information contained in plain text and XML messages.

Table 58-6 Reactive Event Message Format

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Time stamp	Date and time stamp of event in ISO time notation: *YYYY-MM-DDTHH:MM:SS*. **Note** The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time
Message name	Name of message. Specific event names are listed in the "Event Triggers" section on page 58-21.	/mml/header/name
Message type	Specifically "Call Home."	/mml/header/type
Message group	Specifically "reactive."	/mml/header/group
Severity level	Severity level of message (see Table 58-4).	/mml/header/level
Source ID	Product type for routing.	/mml/header/source
Device ID	Unique device identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: *type@Sid@serial*, where • Type is the product model number from backplane SEEPROM. • @ is a separator character. • Sid is "C" identifying the serial ID as a chassis serial number· • Serial is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/ header/deviceId
Customer ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/ header/customerID
Contract ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/ header /contractId
Site ID	Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/ header/siteId
Server ID	If the message is generated from the fabric switch, it is the unique device identifier (UDI) of the switch. Format: *type@Sid@serial*, where • Type is the product model number from backplane SEEPROM. • @ is a separator character. • Sid is "C" identifying the serial ID as a chassis serial number· • Serial is the number identified by the Sid field. Example: "DS-C9509@C@12345678	/mml/header/serverId
Message description	Short text describing the error.	/mml/body/msgDesc
Device name	Node that experienced the event. This is the host name of the device.	/mml/body/sysName
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact
Contact e-mail	E-mail address of person identified as contact for this unit.	/mml/body/sysContactEmail

Table 58-6 Reactive Event Message Format (continued)

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhone Number
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress
Model name	Model name of the switch. This is the specific model as part of a product family name.	/mml/body/chassis/name
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo
Chassis part number	Top assembly number of the chassis.	/mml/body/chassis/partNo
Chassis hardware version	Hardware version of chassis.	/mml/body/chassis/hwVersion
Supervisor module software version	Top level software version.	/mml/body/chassis/swVersion
Affected FRU name	Name of the affected FRU generating the event message.	/mml/body/fru/name
Affected FRU serial number	Serial number of affected FRU.	/mml/body/fru/serialNo
Affected FRU part number	Part number of affected FRU.	/mml/body/fru/partNo
FRU slot	Slot number of FRU generating the event message.	/mml/body/fru/slot
FRU hardware version	Hardware version of affected FRU.	/mml/body/fru/hwVersion
FRU software version	Software version(s) running on affected FRU.	/mml/body/fru/swVersion
Command output name	The exact name of the issued command.	/mml/attachments/attachment/ name
Attachment type	Specifically command output.	/mml/attachments/attachment/ type
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment/ mime
Command output text	Output of command automatically executed (see Table 58-3).	/mml/attachments/attachment/ atdata

Table 58-7 Inventory Event Message Format

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Time stamp	Date and time stamp of event in ISO time notation: *YYYY-MM-DDTHH:MM:SS*. **Note** The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time
Message name	Name of message. Specifically "Inventory Update" Specific event names are listed in the "Event Triggers" section on page 58-21.	/mml/header/name
Message type	Specifically "Inventory Update".	/mml/header/type
Message group	Specifically "proactive".	/mml/header/group
Severity level	Severity level of inventory event is level 2 (seeTable 58-4).	/mml/header/level
Source ID	Product type for routing at Cisco. Specifically "MDS 9000"	/mml/header/source
Device ID	Unique Device Identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: *type@Sid@serial*, where • Type is the product model number from backplane SEEPROM. • @ is a separator character. • Sid is "C" identifying the serial ID as a chassis serial number· • Serial is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/ header /deviceId
Customer ID	Optional user-configurable field used for contact info or other ID by any support service.	/mml/ header /customerID
Contract ID	Optional user-configurable field used for contact info or other ID by any support service.	/mml/ header /contractId
Site ID	Optional user-configurable field, can be used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/ header /siteId
Server ID	If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch. Format: *type@Sid@serial*, where • Type is the product model number from backplane SEEPROM. • @ is a separator character. • Sid is "C" identifying the serial ID as a chassis serial number· • Serial is the number identified by the Sid field. Example: "DS-C9509@C@12345678	/mml/header/serverId
Message description	Short text describing the error.	/mml/body/msgDesc
Device name	Node that experienced the event.	/mml/body/sysName
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact
Contact e-mail	E-mail address of person identified as contact for this unit.	/mml/body/sysContactEmail

Table 58-7 Inventory Event Message Format (continued)

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhone Number
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress
Model name	Model name of the unit. This is the specific model as part of a product family name.	/mml/body/chassis/name
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo
Chassis part number	Top assembly number of the chassis.	/mml/body/chassis/partNo
Chassis hardware version	Hardware version of chassis.	/mml/body/chassis/hwVersion
Supervisor module software version	Top level software version.	/mml/body/chassis/swVersion
FRU name	Name of the affected FRU generating the event message.	/mml/body/fru/name
FRU s/n	Serial number of FRU.	/mml/body/fru/serialNo
FRU part number	Part number of FRU.	/mml/body/fru/partNo
FRU slot	Slot number of FRU.	/mml/body/fru/slot
FRU hardware version	Hardware version of FRU.	/mml/body/fru/hwVersion
FRU software version	Software version(s) running on FRU.	/mml/body/fru/swVersion
Command output name	The exact name of the issued command.	/mml/attachments/attachment /name
Attachment type	Specifically command output.	/mml/attachments/attachment /type
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment /mime
Command output text	Output of command automatically executed after event categories (see "Event Triggers" section on page 58-21).	/mml/attachments/attachment /atdata

Table 58-8 User-Generated Test Message Format

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Time stamp	Date and time stamp of event in ISO time notation: *YYYY-MM-DDTHH:MM:SS*. **Note** The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time
Message name	Name of message. Specifically test message for test type message. Specific event names listed in the "Event Triggers" section on page 58-21).	/mml/header/name
Message type	Specifically "Test Call Home".	/mml/header/type
Message group	This field should be ignored by the receiving Call Home processing application, but may be populated with either "proactive" or "reactive".	/mml/header/group
Severity level	Severity level of message, test Call Home message (see Table 58-4).	/mml/header/level
Source ID	Product type for routing.	/mml/header/source
Device ID	Unique device identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: *type@Sid@serial*, where • Type is the product model number from backplane SEEPROM. • @ is a separator character. • Sid is "C" identifying the serial ID as a chassis serial number· • Serial is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/ header /deviceId
Customer ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/ header /customerId
Contract ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/ header /contractId
Site ID	Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/ header /siteId
Server ID	If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch. Format: *type@Sid@serial*, where • Type is the product model number from backplane SEEPROM. • @ is a separator character. • Sid is "C" identifying the serial ID as a chassis serial number· • Serial is the number identified by the Sid field. Example: "DS-C9509@C@12345678	/mml/header/serverId
Message description	Short text describing the error.	/mml/body/msgDesc
Device name	Switch that experienced the event.	/mml/body/sysName

Table 58-8 User-Generated Test Message Format (continued)

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact
Contact Email	E-mail address of person identified as contact for this unit.	/mml/body/sysContactEmail
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhone Number
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress
Model name	Model name of the switch. This is the specific model as part of a product family name.	/mml/body/chassis/name
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo
Chassis part number	Top assembly number of the chassis. For example, 800-xxx-xxxx.	/mml/body/chassis/partNo
Command output text	Output of command automatically executed after event categories listed in Table 58-3.	/mml/attachments/attachment/atdata
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment/mime
Attachment type	Specifically command output.	/mml/attachments/attachment/type
Command output name	The exact name of the issued command.	/mml/attachments/attachment/name

CHAPTER 59

Configuring Fabric Configuration Servers

This chapter describes the Fabric Configuration Server (FCS) feature provided in the Cisco MDS 9000 Family of directors and switches. It includes the following sections:

About FCS

The Fabric Configuration Server (FCS) provides discovery of topology attributes and maintains a repository of configuration information of fabric elements. A management application is usually connected to the FCS on the switch through an N port. The FCS views the entire fabric based on the following objects:

- Interconnect element (IE) object—Each switch in the fabric corresponds to an IE object. One or more IE objects form a fabric.
- Port object—Each physical port in an IE corresponds to a port object. This includes the switch ports (xE, Fx, and TL ports) and their attached Nx ports.
- Platform object—A set of nodes may be defined as a platform object to make it a single manageable entity. These nodes are end-devices (host systems, storage subsystems) attached to the fabric. Platform objects reside at the edge switches of the fabric.

Each object has its own set of attributes and values. A null value may also be defined for some attributes.

In the Cisco MDS 9000 Family switch environment, multiple VSANs constitute a fabric, where one instance of the FCS is present per VSAN.

If you have attached a management application to a switch, all the frames directed towards the FCS in the switch are part of the port VSAN in the switch port (Fx port). Hence your view of the management application is limited only to this VSAN. However, information about other VSANs that this switch is part of can be obtained either through the SNMP or CLI.

In Figure 59-1 Management Application 1 (M1) is connected through an F port with port VSAN ID 1, and Management Application 2 (M2) is connected through an F port with port VSAN ID 2. M1 can query the FCS information of switches S1 and S3, and M2 can query switches S3 and S4. Switch S2

information is not known to both of them. FCS operations can be done only on those switches that are visible in the VSAN. Note that M2 can send FCS requests only for VSAN 2 even though S3 is also a part of VSAN 1.

Figure 59-1 FCSs in a VSAN Environment

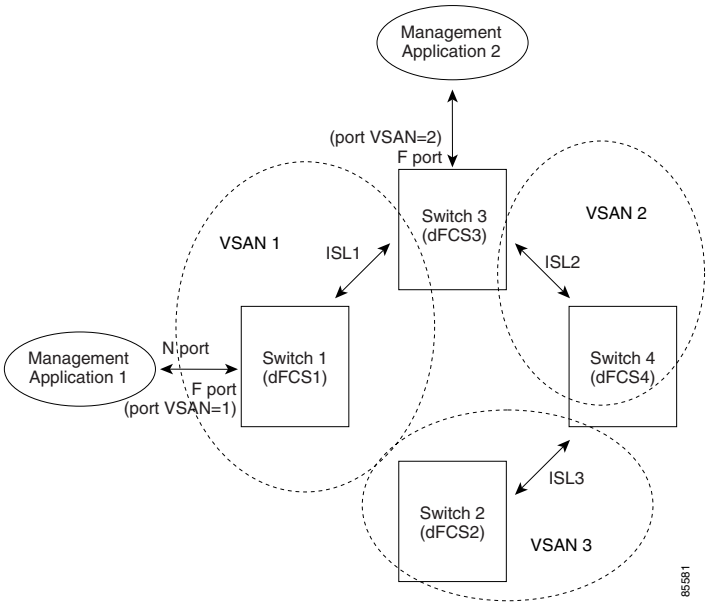

Significance of FCS

This section lists the significance of FCSs.

- FCSs support network management including the following:
  - N port management application can query and obtain information about fabric elements.
  - SNMP manager can use the FCS management information base (MIB) to start discovery and obtain information about the fabric topology.
- FCSs support TE and TL ports in addition to the standard F and E ports.
- FCS can maintain a group of modes with a logical name and management address when a platform registers with it. FCSs maintain a backup of all registrations in secondary storage and update it with every change. When a restart or switchover happens, FCSs retrieve the secondary storage information and rebuild its database.
- SNMP manager can query FCSs for all IEs, ports, and platforms in the fabric.

Displaying FCS Discovery

To display FCS discovery information using Device Manager, follow these steps:

Step 1 Choose **FC > Advanced > Fabric Config Server**.

You see the Fabric Config Server dialog box shown in Figure 59-2.

Figure 59-2 Fabric Config Server Dialog Box

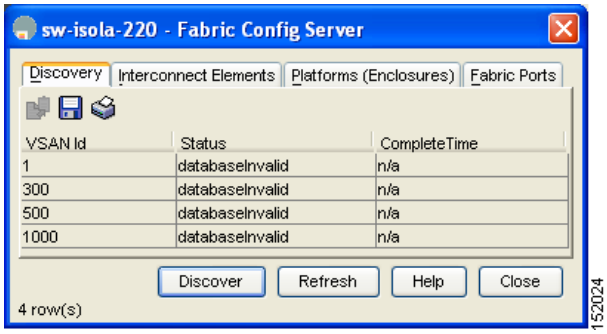

Step 2 Click the **Discovery** tab.

Step 3 Click **Discover** to rediscover the fabric, or click **Refresh** to update the display.

Displaying FCS Elements

To display FCS interconnect element information using Device Manager, follow these steps:

Step 1 Choose **FC > Advanced > Fabric Config Server**.

You see the Fabric Config Server dialog box.

Step 2 Select the **Interconnect Elements** tab.

You see the Fabric Config Server dialog box shown in Figure 59-3.

Figure 59-3 Fabric Config Server Dialog Box

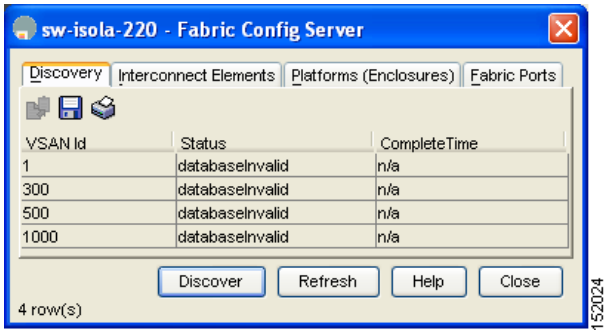

Creating an FCS Platform

To create an FCS platform using Device Manager, follow these steps:

Step 1 Choose **FC > Advanced > Fabric Config Server**.

You see the Fabric Config Server dialog box.

Step 2 Click the **Platforms (Enclosures)** tab.

Step 3 Click **Create**.

You see the Create Fabric Config Server dialog box shown in Figure 59-4.

Figure 59-4 Create Fabric Config Server Dialog Box

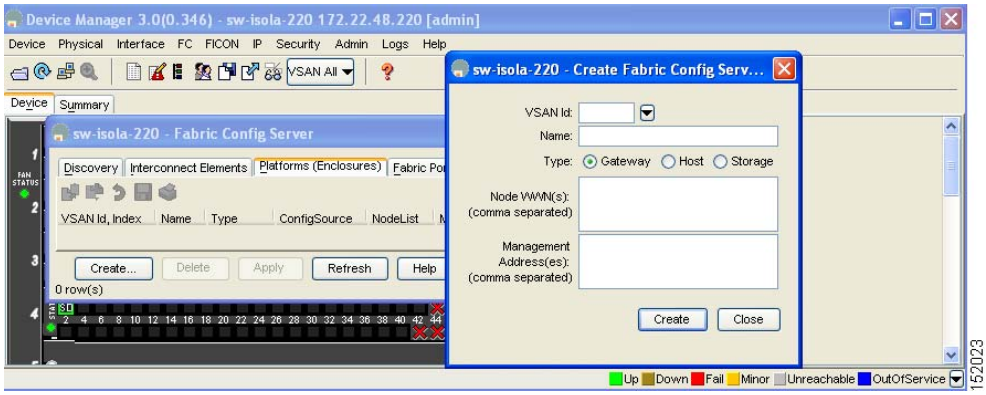

Step 4 Enter the VSAN ID, or select the ID from the drop-down list of available VSAN IDs.

Step 5 Enter the Fabric Configuration Server name in the Name field.

Step 6 Choose the type of server (**Gateway**, **Host**, **Storage**).

Step 7 Enter the WWNs for the server.

Step 8 Enter the management addresses for the server.

Step 9 Click **Create** to create the server, or click **Close** to discard your changes and return to the Fabric Config Server dialog box.

Displaying FCS Fabric Ports

To display FCS discovery information using Device Manager, follow these steps:

Step 1 Choose **FC > Advanced > Fabric Config Server**.

You see the Fabric Config Server dialog box.

Step 2 Click the **Fabric Ports** tab.

You see a list of fabric ports (see Figure 59-5).

Figure 59-5 FCS Fabric Ports

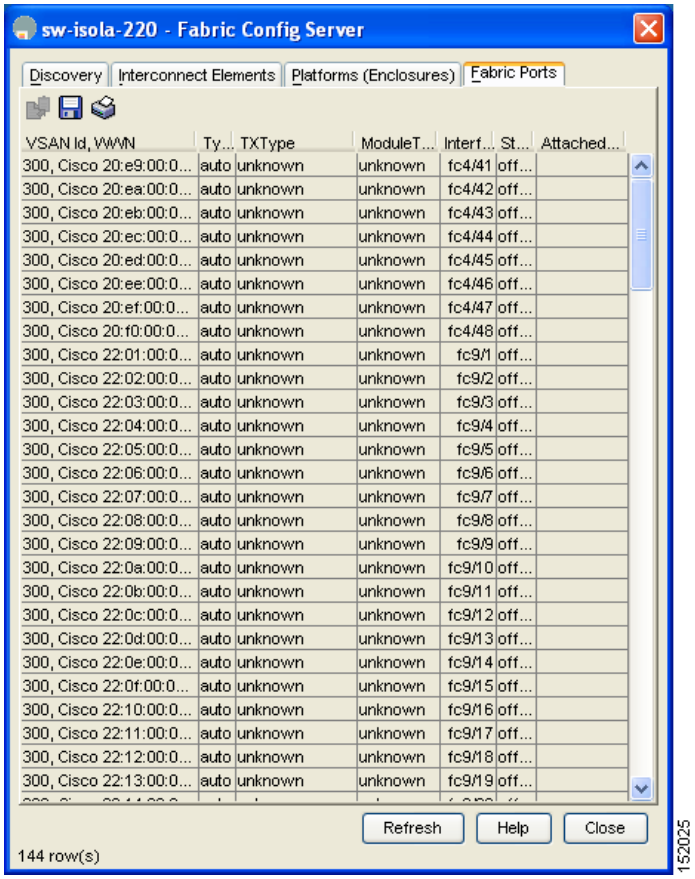

Step 3 Click **Refresh** to update the display.

Default Settings

Table 59-1 lists the default FCS settings.

Table 59-1 Default FCS Settings

Parameters	Default
Global checking of the platform name	Disabled.
Platform node type	Unknown.

Cisco Systems

Send documentation comments to mdsfeedback-doc@cisco.com

PART 9

Traffic Management

CHAPTER 60

Configuring Fabric Congestion Control and QoS

Fibre Channel Congestion Control (FCC) is a Cisco proprietary flow control mechanism that alleviates congestion on Fibre Channel networks.

Quality of service (QoS) offers the following advantages:

- Provides relative bandwidth guarantee to application traffic.
- Controls latency experienced by application traffic.
- Prioritizes one application over another (for example, prioritizing transactional traffic over bulk traffic) through bandwidth and latency differentiation.

This chapter provides details on the QoS and FCC features provided in all switches. It includes the following sections:

About FCC

FCC reduces the congestion in the fabric without interfering with the standard Fibre Channel protocols. The FCC protocol increases the granularity and the scale of congestion control applied to any class of traffic (see Figure 60-1).

Figure 60-1 FCC Mechanisms

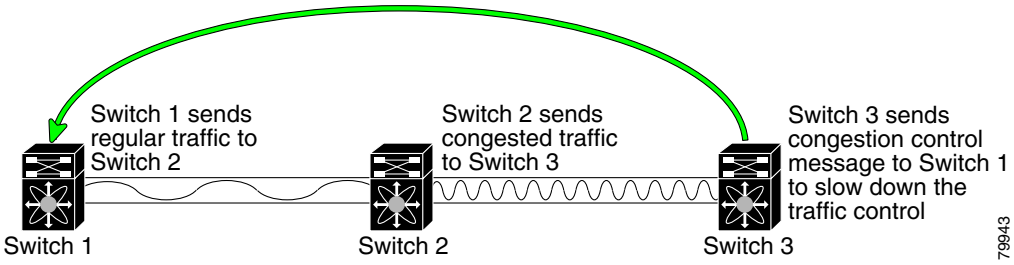

Edge quench congestion control provides feedback to the source about the rate at which frames should be injected into the network (frame intervals).

This section contains the following topics:

FCC Process

When a node in the network detects congestion for an output port, it generates an edge quench message. These frames are identified by the Fibre Channel destination ID (DID) and the source ID. A switch from other vendors simply forwards these frames.

Any receiving switch in the Cisco MDS 9000 Family handles frames in one of these ways:

- It forwards the frame.
- It limits the rate of the frame flow in the congested port.

The behavior of the flow control mechanism differs based on the Fibre Channel DID:

- If the Fibre Channel DID is directly connected to one of the switch ports, the input rate limit is applied to that port.
- If the destination of the edge quench frame is a Cisco domain or the next hop is a Cisco MDS 9000 Family switch, the frame is forwarded.
- If neither of these mechanisms is true, then the frame is processed in the port going towards the FC DID.

All switches (including the edge switch) along the congested path process path quench frames. However, only the edge switch processes edge quench frames.

Enabling FCC

By default, the FCC protocol is disabled. FCC can only be enabled for the entire switch.

Tip If you enable FCC, be sure to enable it in all switches in the fabric.

To enable or disable the FCC feature, using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **FC Services** and then select **FCC** in the Physical Attributes pane.
The FCC information is displayed in the Information pane. The **General** tab is the default.

Step 2 Select the switch on which you want to enable FCC.

Step 3 Check the **Enable** check box.

Step 4 Click **Apply Changes** to save your changes.

Assigning FCC Priority

To assign FCC priority, follow these steps

Step 1 Expand **Switches**, expand **FC Services** and then select **FCC** in the Physical Attributes pane.
The FCC information is displayed in the Information pane. The **General** tab is the default.

Step 2 Select the switch for which you want to assign the FCC priority.

Step 3 Enter the priority in the **Priority** column.

Step 4 Click **Apply Changes** to save your changes..

QoS

QoS implementation in the Cisco MDS 9000 Family follows the differentiated services (DiffServ) model. The DiffServ standard is defined in RFCs 2474 and 2475.

All switches support the following types of traffic:

About Control Traffic

The Cisco MDS 9000 Family supports QoS for internally and externally generated control traffic. Within a switch, control traffic is sourced to the supervisor module and is treated as a high priority frame. A high priority status provides absolute priority over all other traffic and is assigned in the following cases:

- Internally generated time-critical control traffic (mostly Class F frames).
- Externally generated time-critical control traffic entering a switch in the Cisco MDS 9000 Family from a another vendor's switch. High priority frames originating from other vendor switches are marked as high priority as they enter a switch in the Cisco MDS 9000 Family.

Enabling or Disabling Control Traffic

By default, the QoS feature for certain critical control traffic is enabled. These critical control frames are assigned the highest (absolute) priority.

Tip We do not recommend disabling this feature as all critical control traffic is automatically assigned the lowest priority once you issue this command.

To enable or disable the high priority assignment for control traffic, follow these steps:

Step 1 Expand **Switches**, expand **FC Services** and then select **QoS** in the Physical Attributes pane.

The QoS control traffic information is displayed in the Information pane. The **Control** tab is default.

Step 2 Select the switch on which you want to enable or disable control traffic.

Step 3 In the Command column, click the drop-down menu and select **enable** or **disable**.

Step 4 Click **Apply Changes** to save your changes..

About Data Traffic

Online transaction processing (OLTP), which is a low volume, latency sensitive application, requires quick access to requested information. Backup processing application require high bandwidth but are not sensitive to latency. In a network that does not support service differentiation, all traffic is treated identically—they experience similar latency and are allocated similar bandwidths. The QoS feature in Cisco MDS 9000 Family switches provides these guarantees.

Data traffic can be prioritized in distinct levels of service differentiation: low, medium, or high priority. You can apply QoS to ensure that Fibre Channel data traffic for your latency-sensitive applications receive higher priority over throughput-intensive applications such as data warehousing (see Figure 60-2).

Figure 60-2 Prioritizing Data Traffic

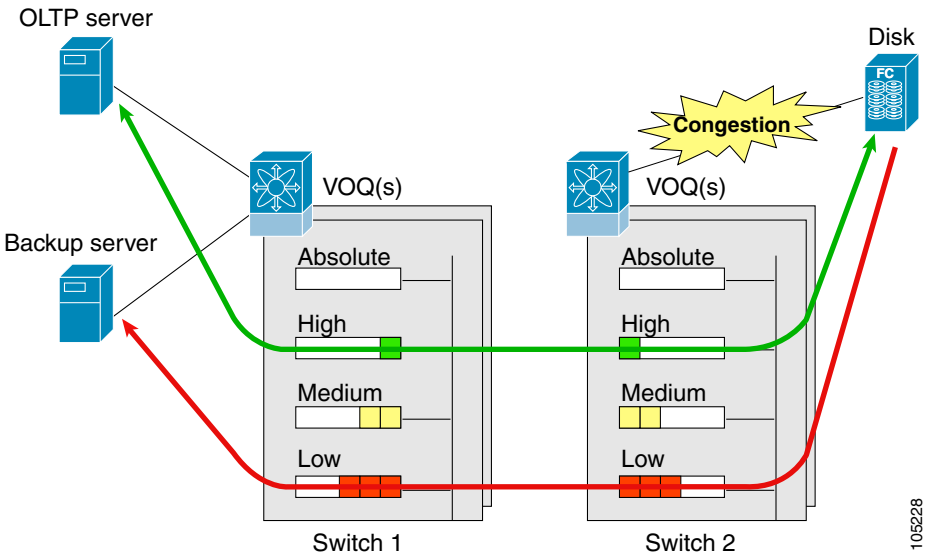

In Figure 60-2, the OLTP traffic arriving at Switch 1 is marked with a high priority level of throughput classification (class map) and marking (policy map). Similarly, the backup traffic is marked with a low priority level. The traffic is sent to the corresponding priority queue within a virtual output queue (VOQ).

A deficit weighted round robin (DWRR) scheduler configured in the first switch ensures that high priority traffic is treated better than low priority traffic. For example, DWRR weights of 70:20:10 implies that the high priority queue is serviced at 7 times the rate of the low priority queue. This guarantees lower delays and higher bandwidths to high priority traffic if congestion sets in. A similar configuration in the second switch ensures the same traffic treatment in the other direction.

If the ISL is congested when the OLTP server sends a request, the request is queued in the high priority queue and is serviced almost immediately as if the high priority queue is not congested. The scheduler assigns its priority over the backup traffic in the low priority queue.

Note When the high priority queue does not have traffic flowing through, the low priority queue uses all the bandwidth and is not restricted to the configured value.

A similar occurrence in Switch 2 sends a response to the transaction request. The round trip delay experienced by the OLTP server is independent of the volume of low priority traffic or the ISL congestion. The backup traffic uses the available ISL bandwidth when it is not used by the OLTP traffic.

Tip To achieve this traffic differentiation, be sure to enable FCC (see the "Enabling FCC" section on page 60-3).

VSAN Versus Zone-Based QoS

While you can configure both zone-based QoS and VSAN-based QoS configurations in the same switch, both configurations have significant differences. Table 60-1 highlights the differences between configuring QoS priorities based on VSANs versus zones.

Table 60-1 QoS Configuration Differences

VSAN-Based QoS	Zone-Based QoS
If you configure the active zone set on a given VSAN and also configure QoS parameters in any of the member zones, you cannot associate the policy map with the VSAN.	You cannot activate a zone set on a VSAN that already has a policy map associated.
If the same flow is present in two class maps associated to a policy map, the QoS value of the class map attached first takes effect.	If the same flow is present in two zones in a given zone set with different QoS values, the higher QoS value is considered.
—	During a zone merge, if the Cisco SAN-OS software detects a mismatch for the QoS parameter, the link is isolated.
Takes effect even if QoS is disabled.	Takes effect only when QoS is enabled.

See the Advanced Zone Attributes, page 26-33 for details on configuring a zone-based QoS policy.

Configuring Data Traffic

To configure QoS using Fabric Manager, follow these steps:

Step 1 Enable the QoS feature.

Step 2 Create and define class maps.

Step 3 Define service policies.

Step 4 Apply the configuration.

About Class Map Creation

Use the class map feature to create and define a traffic class with match criteria to identify traffic belonging to that class. The class map name is restricted to 63 alphanumeric characters and defaults to the match-all option. Flow-based traffic uses one of the following values:

- WWN—The source WWN or the destination WWN.
- Fibre Channel ID (FC ID) —The source ID (SID) or the destination ID (DID). The possible values for mask are FFFFFF (the entire FC ID is used—this is the default), FFFF00 (only domain and area FC ID is used), or FF0000 (only domain FC ID is used).

Note A SID or DID of 0x000000 is not allowed.

- Source interface—The ingress interface.

Tip The order of entries to be matched within a class map is not significant.

Creating a Class Map

To create a class map, using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **FC Services** and then select **QoS** in the Physical Attributes pane.

The QoS information is displayed in the Information pane shown in Figure 60-3. The **Control** tab is default.

Figure 60-3 Quality of Service Control Tab

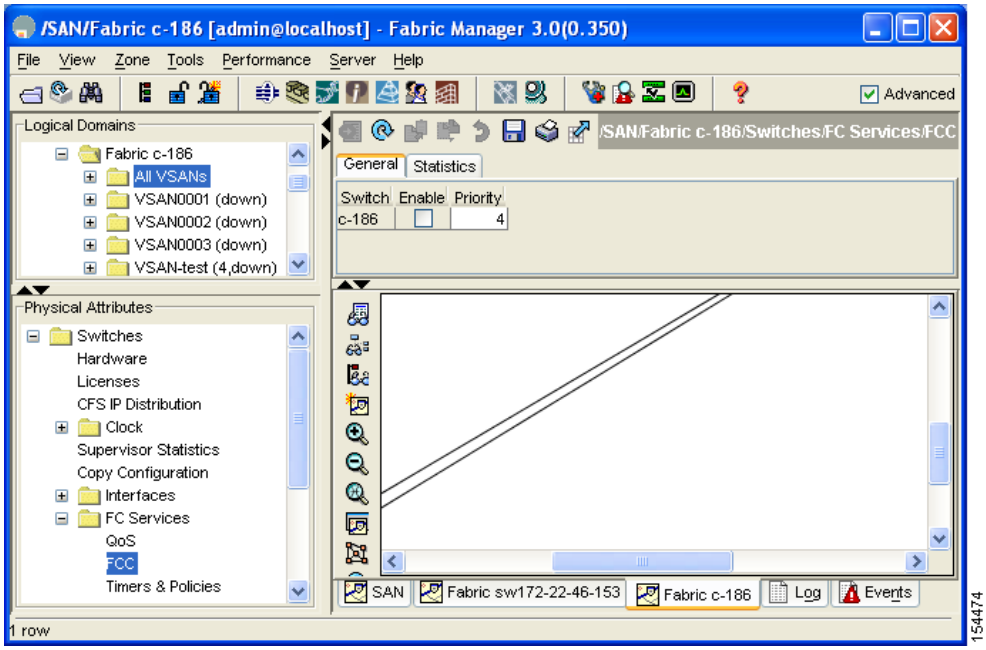

Step 2 In the **Class Maps** tab, click **Create Row** to create a new class map.

You see the Create Class Maps dialog box shown in Figure 60-4.

Figure 60-4 Create Class Maps Dialog Box

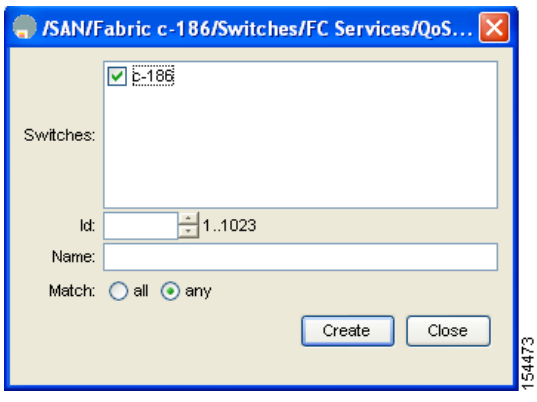

Step 3 Select the switches for the class map (see Figure 60-4).

Step 4 Enter the source ID or the destination ID in the field (see Figure 60-4).

Step 5 Enter a name for the class map (see Figure 60-4).

Step 6 Select a Match mode.You can either match **any** or **all** criterion with one match statement from the class map configuration mode (see Figure 60-4).

Step 7 Click **Create** to proceed with creating the class map or click **Close** to discard any changes (see Figure 60-4).

About Service Policy Definition

Service policies are specified using policy maps. Policy maps provide an ordered mapping of class maps to service levels. You can specify multiple class maps within a policy map, and map a class map to a high, medium, or low service level. The default priority is low. The policy map name is restricted to 63 alphanumeric characters.

As an alternative, you can map a class map to a differentiated services code point (DSCP).The DSCP is an indicator of the service level for a specified frame. The DSCP value ranges from 0 to 63, and the default is 0. A DSCP value of 46 is disallowed.

The order of the class maps within a policy map is important to determine the order in which the frame is compared to class maps. The first matching class map has the corresponding priority marked in the frame.

Note Refer to http://www.cisco.com/warp/public/105/dscpvalues.html#dscpandassuredforwardingclasses for further information on implementing QoS DSCP values.

Note Class maps are processed in the order in which they are configured in each policy map.

About Service Policy Enforcement

When you have configured a QoS data traffic policy, you must enforce the data traffic configuration by applying that policy to the required VSAN(s). If you do not apply the policy to a VSAN, the data traffic configuration is not enforced. You can only apply one policy map to a VSAN.

Note You can apply the same policy to a range of VSANs.

About the DWRR Queue

The Cisco SAN-OS software supports four scheduling queues:

- Strict priority queues are queues that are serviced in preference to other queues—it is always serviced if there is a frame queued in it regardless of the state of the other queues.
- QoS assigns all other traffic to the DWRR scheduling high, medium, and low priority traffic queues.

The DWRR scheduler services the queues in the ratio of the configured weights. Higher weights translate to proportionally higher bandwidth and lower latency. The default weights are 50 for the high queue, 30 for the medium queue, and 20 for the low queue. Decreasing order of queue weights is mandated to ensure the higher priority queues have a higher service level, though the ratio of the configured weights can vary (for example, one can configure 70:30:5 or 60:50:10 but not 50:70:10).

Changing the Weight in a DWRR Queue

To change the weight in a DWRR queue using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **FC Services** and then select **QoS** in the Physical Attributes pane.

The QoS control traffic information is displayed in the Information pane shown in Figure 60-5. The **Control** tab is default.

Figure 60-5 Quality of Service Control Tab

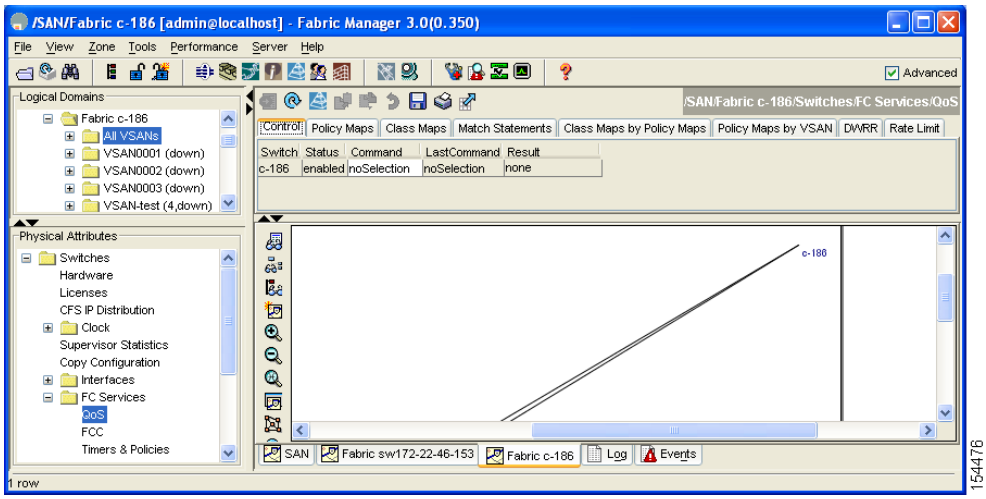

Step 2 Click the **DWWR** tab.

You see the queue status and weight in Figure 60-6.

Figure 60-6 QoS Queue Status and Weight

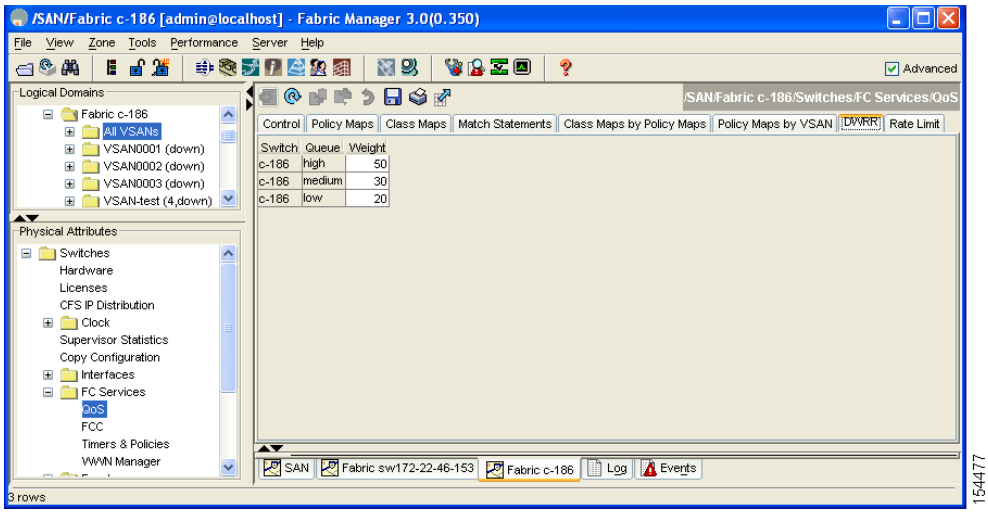

Step 3 Select a switch and change the weight.

Step 4 Click **Apply Changes** to save your changes.

Example Configuration

This section describes a configuration example for the application illustrated in Figure 60-7.

Figure 60-7 Example Application for Traffic Prioritization

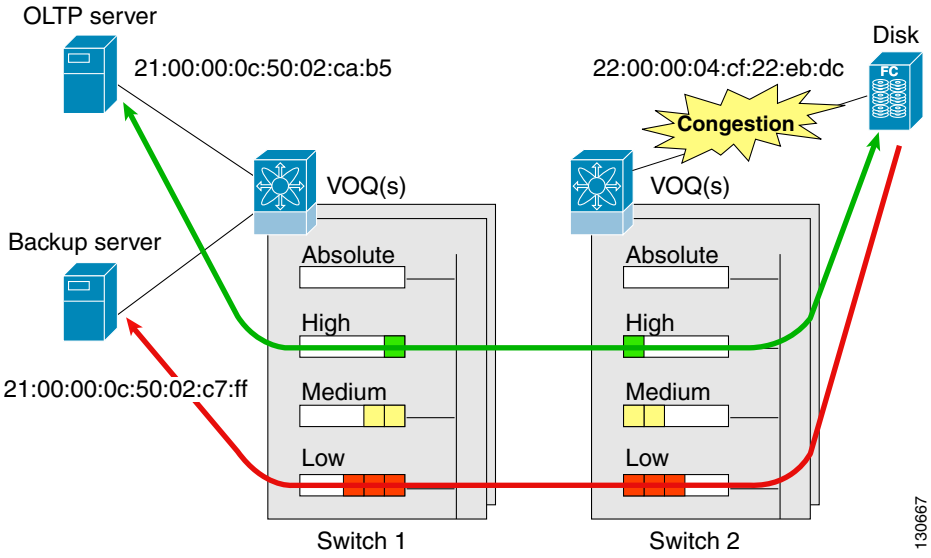

Both the OLTP server and the backup server are accessing the disk. The backup server is writing large amounts of data to the disk. This data does not require specific service guarantees. The volumes of data generated by the OLTP server to the disk are comparatively much lower but this traffic requires faster response because transaction processing is a low latency application.

The point of congestion is the link between Switch 2 and the disk, for traffic from the switch to the disk. The return path is largely uncongested as there is little backup traffic on this path.

Service differentiation is needed at Switch 2 to prioritize the OLTP-server-to-disk traffic higher than the backup-server-to-disk traffic.

To configure traffic prioritization for the example application, follow these steps:

Step 1 Create the class maps.

Step 2 Create the policy map.

Step 3 Assign the service policy.

Step 4 Assign the weights for the DWRR queues.

Step 5 Repeat Step 1 through Step 4 on Switch 1 to address forward path congestion at both switches.

Congestion could occur anywhere in the example configuration. To address congestion of the return path at both switches, you need to create two more class maps and include them in the policy map as follows:

Step 1 Create two more class maps.

Step 2 Assign the class maps to the policy map.

Step 3 Repeat Step 1 through Step 2 on Switch 1 to address return path congestion at both switches.

Ingress Port Rate Limiting

A port rate limiting feature helps control the bandwidth for individual Fibre Channel ports. Port rate limiting is also referred to as ingress rate limiting because it controls ingress traffic into a Fibre Channel port. The feature controls traffic flow by limiting the number of frames that are transmitted out of the exit point on the MAC. Port rate limiting works on all Fibre Channel ports. The rate limit ranges from 1 to 100% and the default is 100%.

Note Port rate limiting can only be configured on Cisco MDS 9100 Series switches, Cisco MDS 9216i switches, and MPS-14/2 modules.

This feature can only be configured if the QoS feature is enabled and if this configuration is performed on a Cisco MDS 9100 series switch, Cisco MDS 9216i switch, or MPS-14/2 module.

To configure the port rate limiting value using Fabric Manager, follow these steps.

Step 1 Expand **Switches**, expand **FC Services** and then select **QoS** in the Physical Attributes pane.

The QoS control traffic information is displayed in the Information pane shown in Figure 60-8. The **Control** tab is default.

Figure 60-8 Quality of Service Control Tab

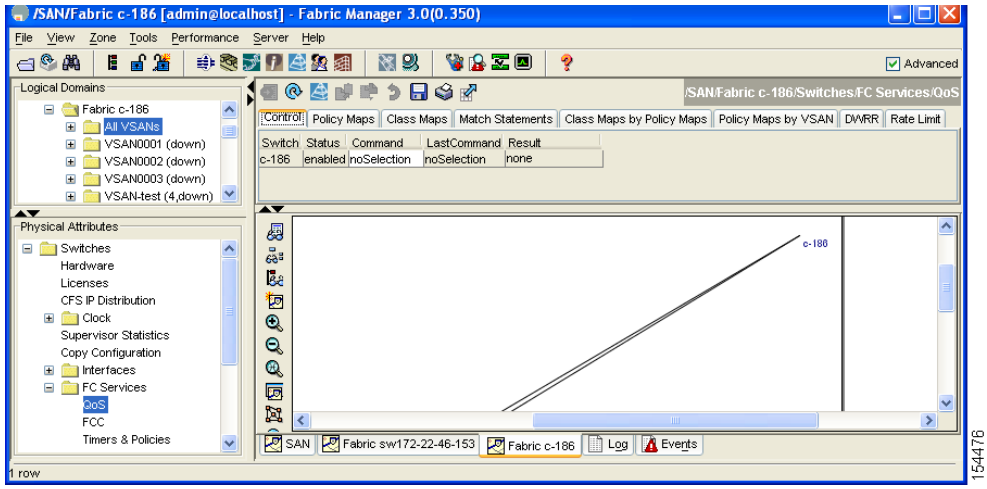

Step 2 Click the **Rate Limit** tab.

You see the information shown in Figure 60-9.

Figure 60-9 Rate Limits for Switch Interfaces

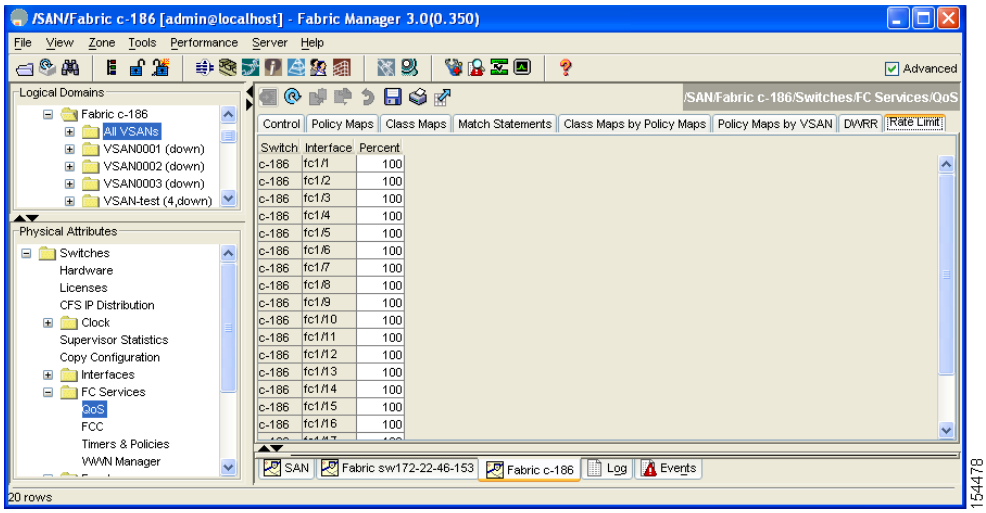

Step 3 Select the switch whose port rate limit you want to change.

Step 4 Enter the desired port rate limit in the Percent column.

Step 5 Click **Apply Changes**.

Default Settings

Table 60-2 lists the default settings for FCC, QoS, and rate limiting features:

Table 60-2 Default FCC, QoS, and Rate Limiting Settings

Parameters	Default
FCC protocol	Disabled.
QoS control traffic	Enabled.
QoS data traffic	Disabled.
Zone-based QoS priority	Low.
Rate limit	100%

CHAPTER 61

Configuring Port Tracking

The Port Tracking feature is unique to the Cisco MDS 9000 Family of switches. This feature uses information about the operational state of the link to initiate a failure in the link that connects the edge device. This process of converting the indirect failure to a direct failure triggers a faster recovery process towards redundant links. When enabled, the port tracking feature brings down the configured links based on the failed link and forces the traffic to be redirected to another redundant link.

This chapter includes the following sections:

About Port Tracking

Generally, hosts can instantly recover from a link failure on a link that is immediately (direct link) connected to a switch. However, recovering from an indirect link failure between switches in a WAN or MAN fabric with a keep-alive mechanism is dependent on several factors such as the time out values (TOVs) and on registered state change notification (RSCN) information (see the "Common Interface Configuration" section on page 32-1 and "About RSCN Information" section on page 29-5).

In Figure 61-1, when the direct link 1 to the host fails, recovery can be immediate. However, when the ISL 2 fails between the two switches, recovery depends on TOVs, RSCNs, and other factors.

Figure 61-1 Traffic Recovery Using Port Tracking

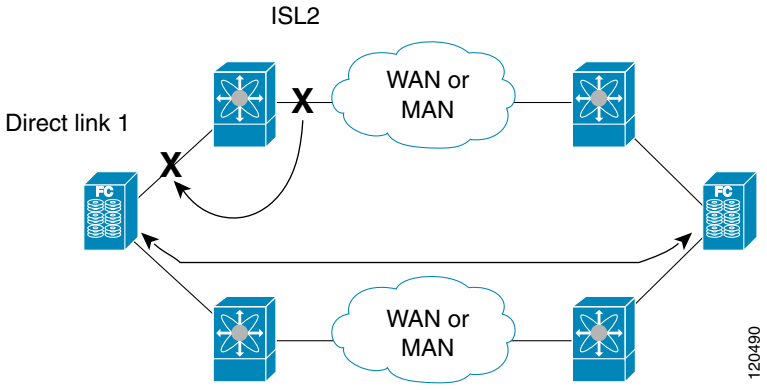

The port tracking feature monitors and detects failures that cause topology changes and brings down the links connecting the attached devices. When you enable this feature and explicitly configure the linked and tracked ports, the Cisco SAN-OS software monitors the tracked ports and alters the operational state of the linked ports on detecting a link state change.

The following terms are used in this chapter.

- Tracked ports—A port whose operational state is continuously monitored. The operational state of the tracked port is used to alter the operational state of one or more ports. Fibre Channel, VSAN, PortChannel, FCIP, or a Gigabit Ethernet port can be tracked. Generally, ports in E and TE port modes can also be Fx ports
- Linked ports—A port whose operational state is altered based on the operational state of the tracked ports. Only a Fibre Channel port can be linked.

Port Tracking

Before configuring port tracking, consider the following guidelines:

- Verify that the tracked ports and the linked ports are on the same Cisco MDS switch.
- Be aware that the linked port is automatically brought down when the tracked port goes down.
- Do not track a linked port back to itself (for example, Port fc1/2 to Port fc2/5 and back to Port fc1/2) to avoid recursive dependency.

This section includes the following topics:

About Port Tracking

Port tracking has the following features:

- The application brings the linked port down when the tracked port goes down. When the tracked port recovers from the failure and comes back up again, the tracked port is also brought up automatically (unless otherwise configured).
- You can forcefully continue to keep the linked port down, even though the tracked port comes back up. In this case, you must explicitly bring the port up when required.

Enabling Port Tracking

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

To configure port tracking, enable the port tracking feature and configure the linked port(s) for the tracked port.

To enable port tracking with Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **Interfaces**, and then select **Port Tracking** in the Physical Attributes pane.

The port tracking information is displayed in the Information pane shown in Figure 61-2. The **Controls** tab is default.

Figure 61-2 Port Tracking

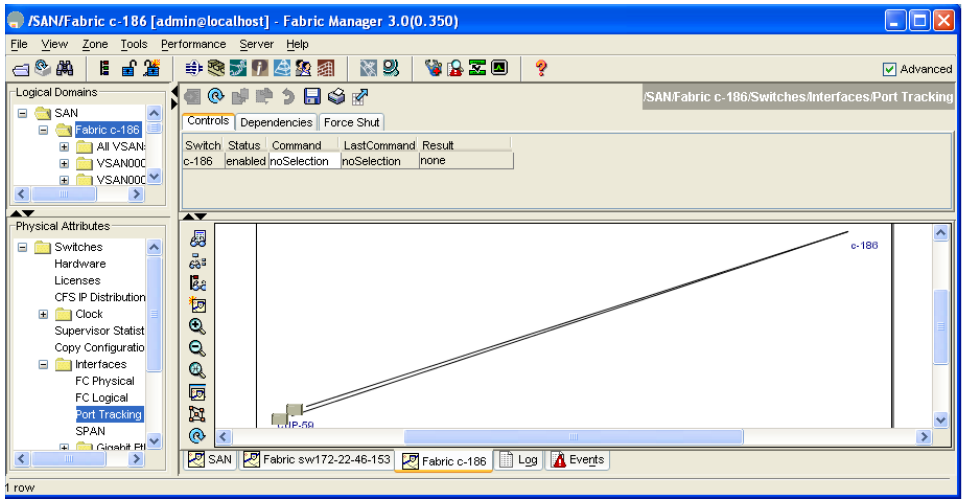

Step 2 Click in the Command column to **enable** or **disable** port tracking.

Depending on your selection the corresponding entry in the Status column changes.

Step 3 Click **Apply Changes** to save your changes.

The entry in the Result column changes to changes to **success**.

About Configuring Linked Ports

You can link ports using one of two methods:

- Operationally binding the linked port(s) to the tracked port (default).
- Continuing to keep the linked port down forcefully—even if the tracked port has recovered from the link failure.

Operationally Binding a Tracked Port

When you configure the first tracked port, operational binding is automatically in effect. When you use this method, you have the option to monitor multiple ports or monitor ports in one VSAN.

To create dependencies, follow these steps:

Step 1 Expand **Switches**, expand **Interfaces**, and then select **Port Tracking** in the Physical Attributes pane.

The port tracking information is displayed in the Information pane. The Controls tab is default.

Figure 61-3 Port Tracking Controls Tab

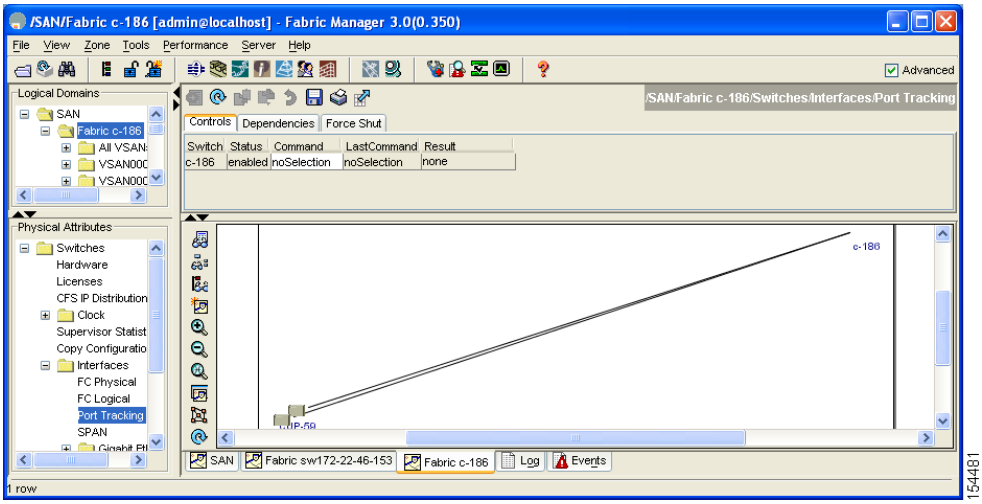

Step 2 Click the **Dependencies** tab.

Step 3 Click **Create Row**.

The Create Port Tracking Dependencies dialog box is displayed (see Figure 61-4).

Figure 61-4 Create Port Tracking Dependencies Dialog Box

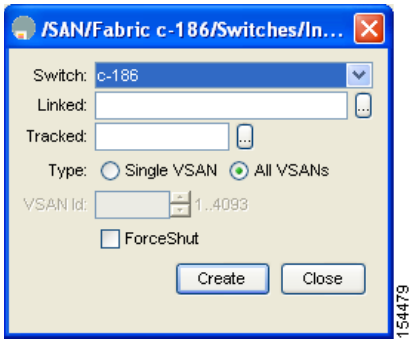

Step 4 Select the switch whose ports you want to track by clicking **...** and selecting from the list.

Step 5 Select the linked port(s) that should be bound to the tracked port(s) by clicking **...** and selecting from the list.

Step 6 Choose **Single VSAN** if you want to track these ports only in one VSAN. Select **All VSANs** if you want to track these ports in all the available VSANs. Also see "About Monitoring Ports in a VSAN" section on page 61-6 for details.

Step 7 If you chose Single VSAN in the previous step, enter the ID of the VSAN where these ports will be monitored.

Step 8 Check the **Forceshut** option if you want to forcefully shutdown the tracked port. Also see "About Forceful Shutdown" section on page 61-6 for details.

Step 9 Click **Create** to proceed with creating this dependency or close the dialog box to discard any changes. If tracking is established, you see **Success** in the lower left corner of the screen shown in Figure 61-5.

Figure 61-5 Successful Port Tracking Established

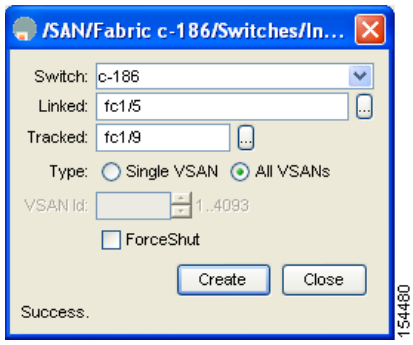

About Tracking Multiple Ports

You can control the operational state of the linked port based on the operational states of multiple tracked ports. When more than one tracked port is associated with a linked port, the operational state of the linked port will be set to down only if all the associated tracked ports are down. Even if one tracked port is up, the linked port will stay up.

In Figure 61-6, only if both ISLs 2 and 3 fail, will the direct link 1 be brought down. Direct link 1 will not be brought down if either 2 or 3 are still functioning as desired.

Figure 61-6 Traffic Recovery Using Port Tracking

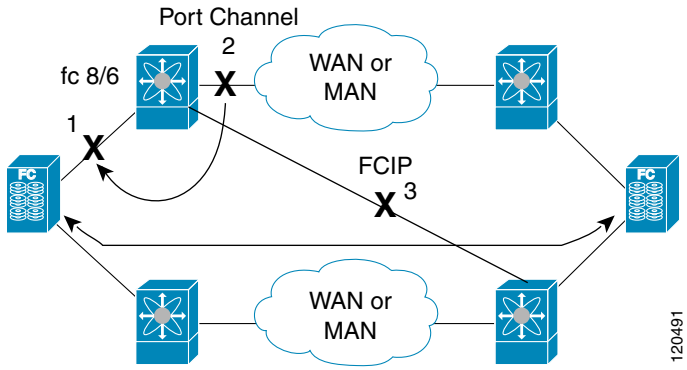

Tracking Multiple Ports

To track multiple ports, see "Operationally Binding a Tracked Port" section on page 61-4.

About Monitoring Ports in a VSAN

You can optionally configure one VSAN from the set of all operational VSANs on the tracked port with the linked port by specifying the required VSAN. This level of flexibility provides higher granularity in tracked ports. In some cases, when a tracked port is a TE port, the set of operational VSANs on the port can change dynamically without bringing down the operational state of the port. In such cases, the port VSAN of the linked port can be monitored on the set of operational VSANs on the tracked port.

If you configure this feature, the linked port is up only when the VSAN is up on the tracked port.

Tip The specified VSAN does not have to be the same as the port VSAN of the linked port.

Monitoring Ports in a VSAN

To monitor a tracked port in a specific VSAN, see "Operationally Binding a Tracked Port" section on page 61-4.

About Forceful Shutdown

If a tracked port flaps frequently, then tacking ports using the operational binding feature may cause frequent topology change. In this case, you may choose to keep the port in the down state until you are able to resolve the reason for these frequent flaps. Keeping the flapping port in the down state forces the traffic to flow through the redundant path until the primary tracked port problems are resolved. When the problems are resolved and the tracked port is back up, you can explicitly enable the interface.

Tip If you configure this feature, the linked port continues to remain in the shutdown state even after the tracked port comes back up. You must explicitly remove the forced shut state (by administratively bringing up this interface) of the linked port once the tracked port is up and stable.

Forcefully Shutdown a Tracked Port

To forcefully shutdown a tracked port, see "Operationally Binding a Tracked Port" section on page 61-4.

Default Settings

Table 61-1 lists the default settings for port tracking parameters.

Table 61-1 Default Port Tracking Parameters

Parameters	Default
Port tracking	Disabled
Operational binding	Enabled along with port tracking

Cisco Systems

Send documentation comments to mdsfeedback-doc@cisco.com

PART 10

Troubleshooting

CHAPTER 62

Troubleshooting Your Fabric

This chapter describes basic troubleshooting methods used to resolve issues with switches. This chapter contains the following sections:

Troubleshooting Tools and Techniques

Multiple techniques and tools are available to monitor and trouble shoot the Cisco MDS 9000 Family of switches. These tools provide a complete, integrated, multi-level analysis solution.

Fabric Manager Server—The Cisco Fabric Manager Server provides a long-term, high level view of storage network performance. Fabric wide performance trends can be analyzed using Performance Manager. It provides the starting point for deeper analysis to resolve network hot-spots.

Device Manager—If a performance problem is detected with the Fabric Manager Server, use Cisco Device Manager to view port level statistics in real-time. Details on protocols, errors, discards, byte and frame counts are available. Samples can be taken as frequently as every 2 seconds, and values can be viewed in text form or graphically as pie, bar, area and line changes.

Traffic Analyzer—Another option is to launch the Cisco Traffic Analyze for Fibre Channel from the Fabric Manager Server to analyze the traffic in greater depth. The Cisco Traffic Analyzer allows you to breakdown traffic by VSANs and protocols and to examine SCSI traffic at a logical unit number (LUN) level.

Protocol Analyzer—If even deeper investigation is needed, the Cisco Protocol Analyzer for Fibre Channel can be launched in-context from the Cisco Traffic Analyzer. The Cisco Protocol Analyzer enables you to examine actual sequences of Fibre Channel frames easily using the Fibre Channel and SCSI decoders Cisco developed for Ethereal.

Port Analyzer Adapter—Fabric Manager Server and Device Manager use SNMP to gather statistics. They fully utilize the built in MDS statistics counters. Even so, there are limits to what the counters can collect.

Integration with the Cisco Traffic Analyzer and Cisco Protocol Analyzer extend the MDS analysis capabilities by analyzing the Fibre Channel traffic itself. The Cisco MDS 9000 Family Switched Port Analyzer (SPAN) enables these solutions via a flexible, non-intrusive technique to mirror traffic selectively from one or more ports to another MDS port within a fabric.

The Cisco Port Analyzer Adapter (PAA) encapsulates SPAN traffic in an Ethernet header for transport to a PC or workstation for analysis. Both Fibre Channel control and data plane traffic are available using SPAN. The PAA broadcasts the Ethernet packets, so they cannot be routed across IP networks. Hubs and switches can be used, provided they are in the same Ethernet subnet. Direct connections between a PAA and the PC are also supported. The PAA can reduce Ethernet traffic by truncating Fibre Channel data.

Both the Cisco Traffic Analyzer and Cisco Protocol Analyzer require the PAA to transport MDS SPAN traffic to a PC or workstation.

Note The Cisco Traffic Analyzer works best with the Cisco Port Analyzer Adapter 2, because it provides a length value for truncated data, enabling accurate byte count reporting.

Cisco Traffic Analyzer

The Cisco Traffic Analyzer for Fibre Channel provides real-time analysis of SPAN traffic or traffic captured previously using the Cisco Protocol Analyzer. The Fibre Channel traffic from multiple Cisco Port Analyzer Adapters (PAA) can be aggregated and analyzed by the Cisco Traffic Analyzer.

There are limits to how may SPAN sources can be sent to a single SPAN destination port on an MDS. Aggregation extends the amount of information that can be analyzed in a unified set of reports by the Cisco Traffic Analyzer.

Note The aggregation capabilities are restricted to the information collect by Ethernet connections to a single PC. Aggregation across multiple PCs is NOT available.

The Cisco Traffic Analyzer presents it reports through a Web server, so you can view them locally or remotely. The traffic analysis functions are provided by ‘ntop’ open-source software, which was enhanced by Cisco to add Fibre Channel and SCSI analysis and MDS enhanced inter-switch link (ISL) header support for SPAN. ntop is available on the Cisco.com software download center, under the Cisco Port Analyzer Adapter. ntop is also available on the Internet at http://www.ntop.org/ntop.html. The Cisco enhanced ntop runs under Microsoft Windows and Linux operating systems.

The Cisco Traffic Analyzer for Fibre Channel presents reports with network wide statistics. The Summary Traffic report shows what percentage of traffic was within different ranges of frames sizes. A breakdown of the percentage of traffic for each protocol like SCSI, ELS, etc. is provided. The average and peak throughput for the SPAN traffic being analyzed are also provided.

Fibre Channel traffic can be analyzed on a per VSAN basis with the Cisco Traffic Analyzer. The Domain Traffic Distribution graphs indicate how much traffic (bytes) were transmitted or received by a switch for a particular VSAN. FC Traffic Matrix graphs show how much traffic is transmitted and received between Fibre Channel sources and destinations. The total byte and frame counts for each VSAN are also provided.

Statistics can be analyzed for individual host and storage ports. You can see the percentage of SCSI read vs. write traffic, SCSI vs. other traffic, and percentage of transmitted vs. received bytes and frames. The peak and average throughput values are available for data transmitted and received by each port.

Cisco Protocol Analyzer

The Cisco Protocol Analyzer for Fibre Channel enables you to view Fibre Channel traffic frames in real-time or from a capture file. Fibre Channel and SCSI decoders enable you to view and analyze traffic at the frame level. It matches response with request for complete decoding, which greatly simplifies navigation. Response time between response and status are presented.

The Cisco Protocol Analyzer is VSAN aware, so VSANs can be used as criteria for capture and display filters, and to colorize the display. VSAN #s can also be displayed in a column. Summary statistics are available for protocol distribution percentages and total bytes/frames transferred between specific Fibre Channel source/destination pairs. File capture and filtering controls are available. Captured files can be analyzed by either the Cisco Protocol Analyzer or the Cisco Traffic Analyzer.

Numerous features have been included for ease-of-use. You can find frames that meet particular criteria and mark them. Entries in the frame (packet) list can be colorized to highlight items of interest, and columns can be added/removed as desired.

The protocol analysis functions are provided by ‘Ethereal’ open-source software, which was enhanced by Cisco to decode Fibre Channel and SCSI protocols and support MDS enhanced inter-switch link (ISL) headers for SPAN. Ethereal is available on the Cisco.com software download center, under the Cisco Port Analyzer Adapter. Ethereal is also available on the Internet at http://www.ethereal.com. Ethereal runs under Microsoft Windows, Solaris, and Linux operating systems.

Analyzing Switch Device Health

The Switch Health option lets you determine the status of the components of a specific switch.

To use the Switch Health option in Fabric Manager to determine the status of the components of a specific switch, follow these steps:

Step 1 Select **Switch Health** from the Tools menu.

You see the Switch Health Analysis window.

Step 2 Click **Start** to identify problems currently affecting the selected switch.

You see any problems listed in the switch health analysis window shown in Figure 62-1.

Figure 62-1 Results of a Switch Health Analysis

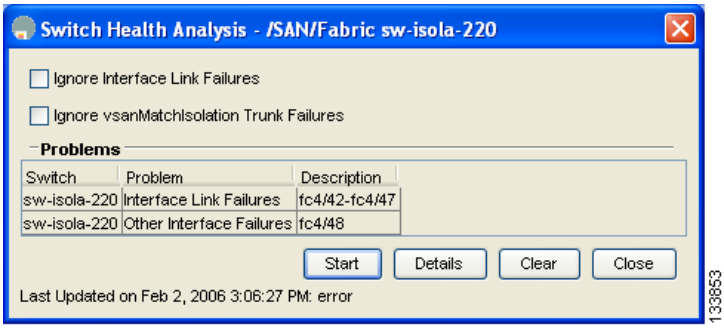

Step 3 Click **Clear** to remove the contents of the Switch Health Analysis window.

Step 4 Click **Close** to close the window.

Analyzing Switch Fabric Configuration

The Fabric Configuration option lets you analyze the configuration of a switch by comparing the current configuration to a specific switch or to a policy file. You can save a switch configuration to a file and then compare all switches against the configuration in the file.

To use the Fabric Configuration option in Fabric Manager to analyze the configuration of a switch, follow these steps:

Step 1 Click **Fabric Configuration** from the Tools menu.

You see the Fabric Configuration Analysis dialog box.

Step 2 Decide whether you want to compare the selected switch to another switch, or to a policy file.

- If you are making a switch comparison, select **Policy Switch** and then click the drop-down arrow to see a list of switches.
- If you are making a policy comparison, select **Policy File**. Then click the **...** button to the right of this option to browse your file system and select a policy file (*.XML).

Step 3 Click **Rules** to set the rules to apply when running the Fabric Configuration Analysis tool.

You see the Rules window.

Step 4 Change the rules as needed and click **OK**.

Step 5 Click **Compare**.

The system analyzes the configuration and displays issues that arise as a result of the comparison as shown in Figure 62-2.

Figure 62-2 Results of a Fabric Configuration Analysis

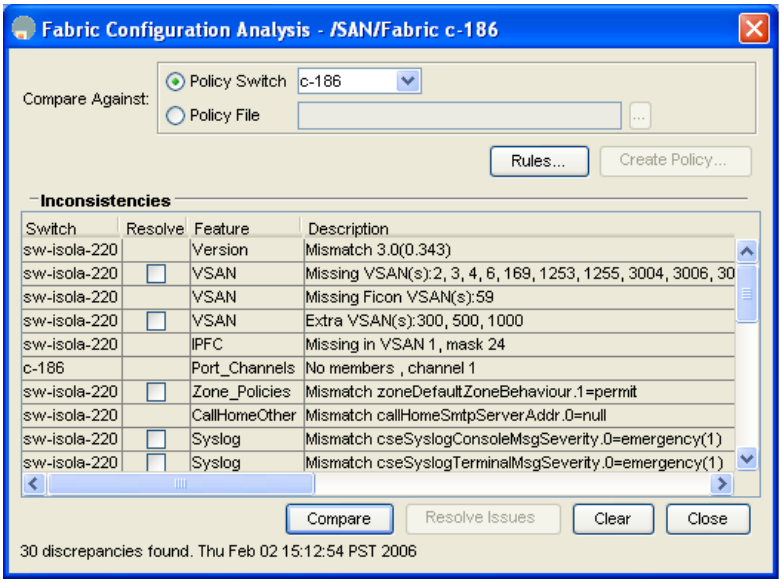

Step 6 Click the check boxes in the Resolve column for the issues you want to resolve.

Step 7 Resolve them by clicking **Resolve Issues**.

Step 8 Click **Clear** to remove the contents of the window.

Step 9 Click **Close** to close the window.

Analyzing End-to-End Connectivity

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

- End to End Connectivity

The ping and redundancy tests are now mutually exclusive, you cannot run both at the same time.

To use the End to End Connectivity option in Fabric Manager to determine connectivity and routes, follow these steps:

Step 1 Choose **End to End Connectivity** from the Tools menu.

You see the End to End Connectivity Analysis dialog box.

Step 2 Select the VSAN whose connectivity will be verified from the VSAN drop-down list.

Step 3 Select whether to perform the analysis for all active zones or for the default zone.

Step 4 Click **Ensure that members can communicate** to perform a Fibre Channel ping between the selected endpoints.

Step 5 Identify the number of packets, the size of each packet, and the time out in milliseconds.

Step 6 Analyze the redundant paths between endpoints by checking the **Ensure that redundant paths exist between members** check box.

Step 7 Check the **Report errors for** check box to see a report of zone and device errors.

Step 8 Click **Analyze**.

The End to End Connectivity Analysis window displays the selected endpoints including the switch to which each is attached, and the source and target ports used to connect it, as shown in Figure 62-3.

Figure 62-3 Results of an End-to-End Connectivity Analysis

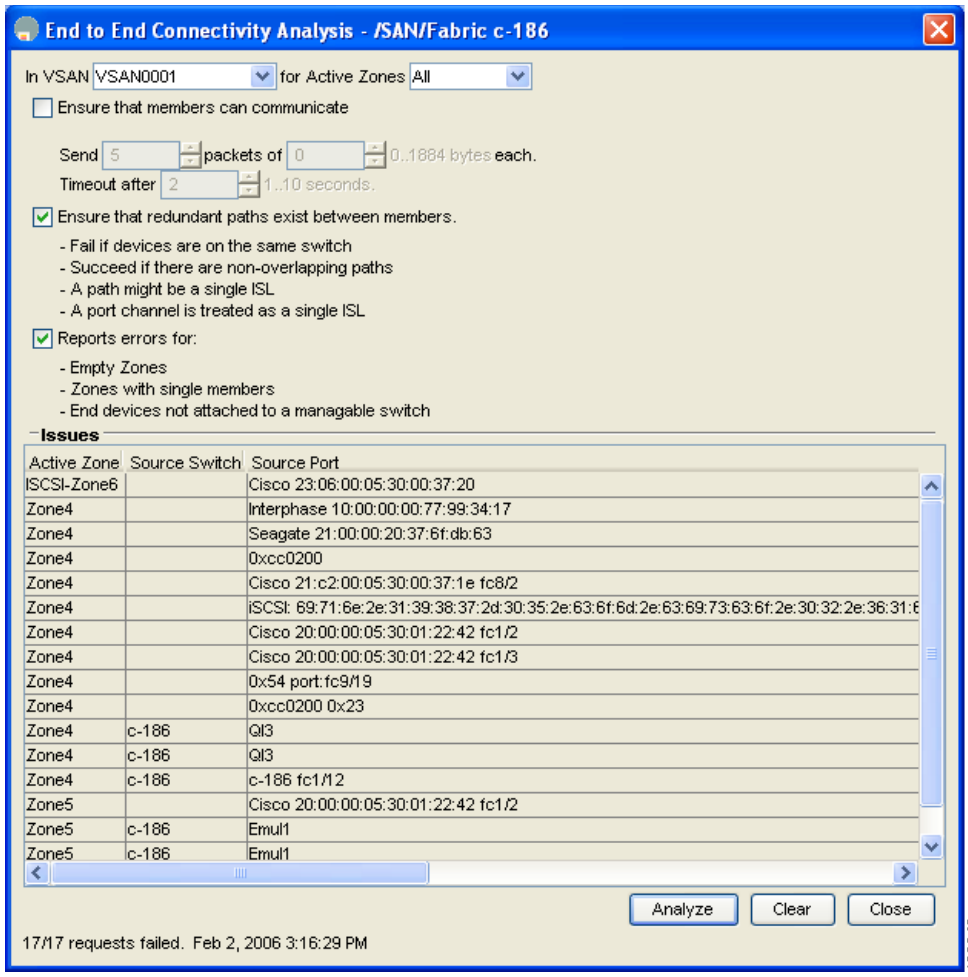

The output shows all the requests that have failed. The possible descriptions are:

- Ignoring empty zone—No requests are issued for this zone.
- Ignoring zone with single member—No requests are issued for this zone.
- Source/Target are unknown—No name server entries exist for the ports or we have not discovered the port during discovery.
- Both devices are on the same switch.
- No paths exist between the two devices.
- VSAN does not have an active zone set and the default zone is denied.
- Average time micro secs—The latency value was more than the threshold supplied.

Step 9 Click **Clear** to remove the contents of the window.

Step 10 Click **Close** to close the window.

Using the Ping Tool (fcping)

You can use the Ping tool to determine connectivity from another switch to a port on your switch.

To use the Ping tool in Fabric Manager to determine connectivity, follow these steps:

Step 1 Select **Ping** from the Tools menu. You can also select it from the right-click context menus for hosts and storage devices in the Fabric pane.

You see the Ping dialog box.

Step 2 Select the source switch from the Source Switch drop-down list.

Step 3 Select the VSAN in which you want to verify connectivity from the VSAN drop-down list.

Step 4 Select the target end port for which to verify connectivity from the Target Endport drop-down list.

Step 5 Click **Start** to perform the ping between your switch and the selected port.

You see the results in the dialog box shown in Figure 62-4.

Figure 62-4 Ping Results

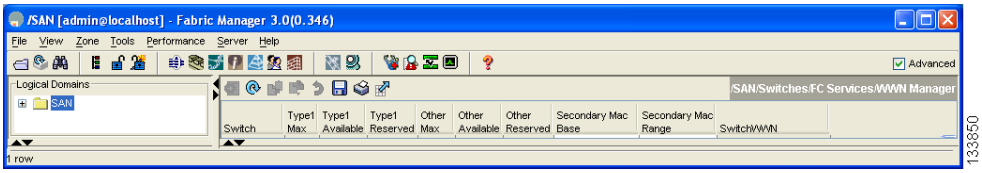

Step 6 Click **Clear** to clear the contents of the window and perform another ping, or click **Close** to close the window.

Using Traceroute (fctrace) and Other Troubleshooting Tools

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

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

To use the Traceroute option in Fabric Manager to verify connectivity, follow these steps:

Step 1 Select **Trace Route** from the Tools menu.

You see the Trace Route dialog box.

Step 2 Select the source switch from the Source Switch drop-down list.

Step 3 Select the VSAN for which to verify connectivity from the VSAN drop-down list.

Step 4 Select the target end port for which to verify connectivity from the Target Endport drop-down list.

Step 5 Click **Start** to perform the traceroute between your switch and the selected port.

You see the results at the bottom of the dialog box as shown in Figure 62-5.

Figure 62-5 Successful Trace Route Results

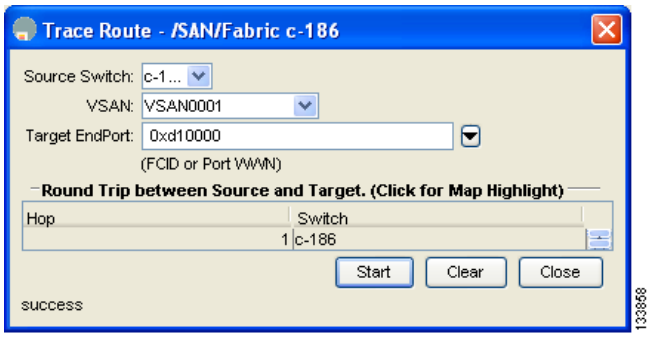

Step 6 Click **Clear** to clear the contents of the window and perform another traceroute, or click **Close** to close the window.

Analyzing the Results of Merging Zones

You can use the Zone Merge option on the Zone menu to determine if two connected switches have compatible zone configurations.

To use the Zone Merge option in Fabric Manager to determine zone configuration compatibility, follow these steps:

Step 1 Choose **Merge Analysis** from the Zone menu.

You see the Zone Merge Analysis dialog box.

Step 2 Select a switch from each drop-down list.

Step 3 Select the VSAN for which you want to perform the zone merge analysis.

Step 4 Repeat Step 3 as needed.

Step 5 Click **Analyze**.

The Zone Merge Analysis window displays any inconsistencies between the zone configuration of the two selected switches as shown in Figure 62-6 .

Figure 62-6 Results of Zone Merge Analysis

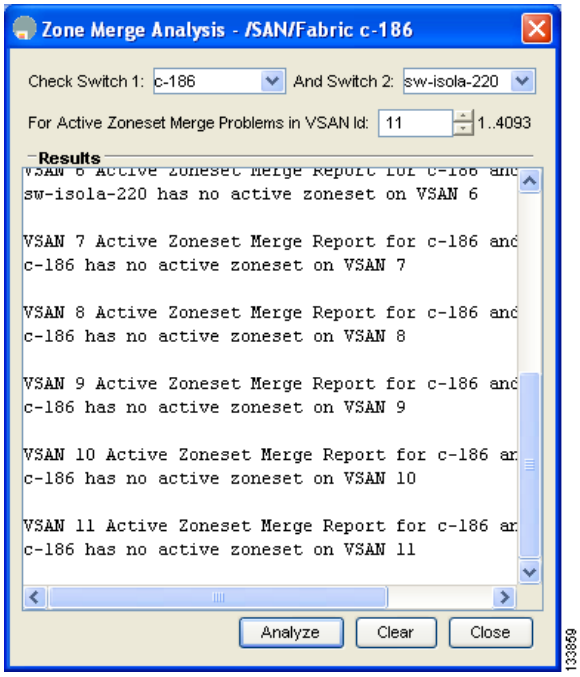

Step 6 Click **Clear** to remove the contents of the window.

Step 7 Click **Close** to close the window.

Issuing the Show Tech Support Command

The **show tech support** command is useful when collecting a large amount of information about your switch for troubleshooting purposes. The output can be provided to technical support representatives when reporting a problem.

You can issue a **show tech support** command from Fabric Manager for one or more switches in a fabric. The results of each command are written to a text file, one file per switch, in a directory you specify. You can then view these files using Fabric Manager.

You can also save the Fabric Manager map as a JPG file. The file is saved with the name of the seed switch (for example, 172.22.94.250.jpg).

You can zip up all the files (the show tech support output and the map file image) and send the resulting zipped file to technical support.

To issue the **show tech support** command using Fabric Manager, follow these steps:

Step 1 Select **Show Tech Support** from the Tools menu.

You see the Show Tech Support dialog box.

Step 2 Select the switches for which to view tech support information by checking the check boxes next to their IP addresses.

Step 3 Set the time-out value.

The default is 30 seconds.

Step 4 Select the folder where you want the text files (containing the tech support information) to be written.

Step 5 Check the **Save Map** check box if you want to save a screenshot of your map as a JPG file.

Step 6 Check the **Compress all files as** check box to compress the files into a zip file.

Step 7 Click **OK** to start issuing the **show tech support** command to the switches you specified, or click **Close** to close the Show Tech Support dialog box without issuing the **show tech support** command (see Figure 62-7).

In the Status column next to each switch, you see a highlighted status. A yellow highlight indicates that the **show tech support** command is currently running on that switch. A red highlight indicates an error. A green highlight like the one shown in Figure 62-7 indicates that the **show tech support** command has completed successfully.

Figure 62-7 Successful Results of the Show Tech Support Command

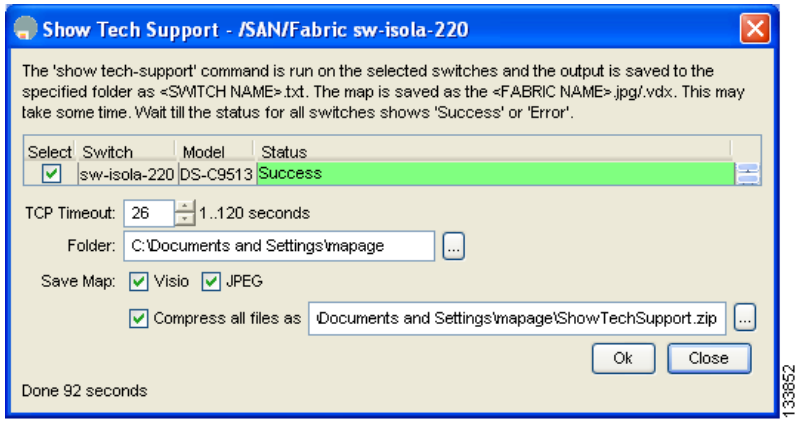

Step 8 If prompted, enter your user name and password in the appropriate fields for the switch in question.

Note that for Fabric Manager to successfully issue the **show tech support** command on a switch, that switch must have this user name and password. Fabric Manager is unable to log into a switch that does not have a user name and password and an error is returned for that switch.

Note If you would like to view output files of the **show tech support** command without using Fabric Manager, open them with any text editor. Each file is named with the switch's IP address and has a .TXT extension (for example, 111.22.33.444.txt).

Locating Other Switches

The Locate Switches option uses SNMPv2 and discovers devices responding to SNMP requests with the read-only community string public. You can use this feature if:

- You have third-party switches that do not implement the FC-GS3 FCS standard that provides management IP addresses.
- You want to locate other Cisco MDS 9000 switches in the subnet but are not physically connected to the fabric (and therefore cannot be found via neighbors).

To locate switches that are not included in the currently discovered fabric using Fabric Manager, follow these steps:

Step 1 Select **Locate Switches and Devices** from the File menu.

You see the Locate Switches dialog box.

Step 2 In the Comma Separated Subnets field, enter a range of specific addresses belonging to a specific subnet to limit the research for the switches. To look for a Cisco MDS 9000 switch belonging to subnet 192.168.199.0, use the following string:

192.168.100.[1-254]

Multiple ranges can be specified, separated by commas. For example, to look for all the devices in the two subnets 192.168.199.0 and 192.169.100.0, use the following string:

192.168.100.[1-254], 192.169.100.[1-254]

Step 3 Enter the appropriate read community string in the Read Community field.

The default value for this string is **public**.

Step 4 Click **Display Cisco MDS 9000 Only** to display only the Cisco MDS 9000 Family switches in your network fabric.

Step 5 Click **Search** to discover switches and devices in your network fabric.

You see the results of the discovery in the Locate Switches window. (See Figure 62-8.)

Figure 62-8 Search Results for Switches and Devices

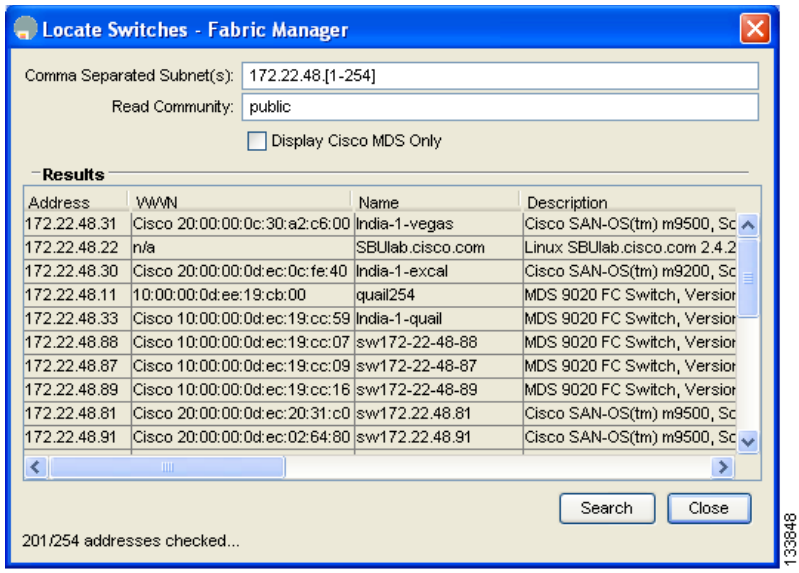

Note The number in the lower left corner of the screen increments as the device locator attempts to discover the devices in your network fabric. When the discovery process is complete, the number indicates the number of rows displayed.

Step 6 Click **Close** to close this dialog box.

Getting Oversubscription Information in Device Manager

To determine oversubscription for a module using Device Manager, follow these steps:

Step 1 Right-click the module you want to check for oversubscription and select **Check Oversubscription** from the pop-up menu.You see the oversubscription dialog box shown in Figure 62-9.

Figure 62-9 Results of Oversubscription Check

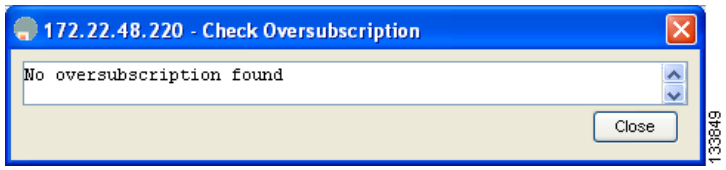

Step 2 Click **Close** to close this dialog box.

Note The module must be capable of oversubscription in order for you to see this menu item.

Fibre Channel Time Out Values

You can modify Fibre Channel protocol related timer values for the switch by configuring the following time out values (TOVs):

- Distributed services TOV (D_S_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 5,000 milliseconds.
- Error detect TOV (E_D_TOV)—The valid range is from 1,000 to 10,000 milliseconds. The default is 2,000 milliseconds. This value is matched with the other end during port initialization.
- Resource allocation TOV (R_A_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 10,000 milliseconds. This value is matched with the other end during port initialization.

Note The fabric stability TOV (F_S_TOV) constant cannot be configured.

Timer Configuration Across All VSANs

You can modify Fibre Channel protocol related timer values for the switch.

Caution The D_S_TOV, E_D_TOV, and R_A_ TOV values cannot be globally changed unless all VSANs in the switch are suspended.

To configure timeouts using Fabric Manager, follow these steps:

Step 1 Select **SAN** in the Logical Domains pane to include all VSANs.

Step 2 Expand **Switches,** expand **FC Services** and select **Timers & Policies** in the Physical Attributes pane.

You see the timers for switches in the Information pane.

Step 3 Click **Change Timeouts** to configure the time-out values.

You see the Change Timeouts dialog box shown in Figure 62-10.

Figure 62-10 Change Timeouts Dialog Box

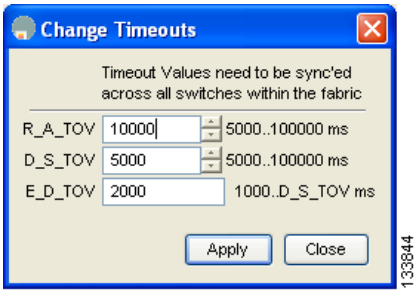

Step 4 Indicate values for R_A_TOV (Resource Allocation Time Out Value), D_S_TOV (Distributed Services Time Out Value), and E_D_TOV (Error Detect Time Out Value).

Step 5 Click **Apply**.

Step 6 Click **Close** to close the dialog box.

To configure timer policies in Device Manager, follow these steps:

Step 1 Choose **FC > Advanced > Timers/Policies**. You see timer policies for a single switch in the dialog box shown in Figure 62-11.

Figure 62-11 Configure Timer Policies in Device Manager

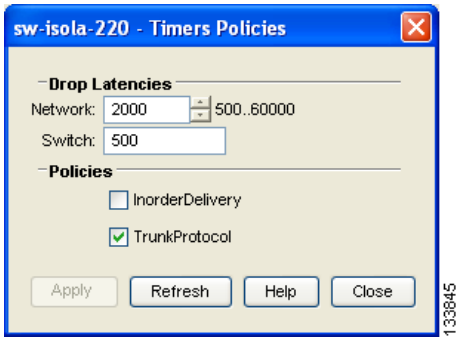

Step 2 Select a network from the drop-down list and specify a switch.

Step 3 Check the check boxes for **InOrderDeliver** and/or **Trunk Protocol.**

Step 4 Click **Apply.**

Step 5 Click **Close** to close the dialog box.

Timer Configuration Per-VSAN

You can also issue an fctimer for a specified VSAN to configure different TOV values for VSANs with special links like FC or IP tunnels. You can configure different E_D_TOV, R_A_TOV, and D_S_TOV values for individual VSANs. Active VSANs are suspended and activated when their timer values are changed.

Caution You cannot perform a nondisruptive downgrade to any earlier version that does not support per-VSAN FC timers.

Note This configuration must be propagated to all switches in the fabric—be sure to configure the same value in all switches in the fabric.

If a switch is downgraded to Cisco MDS SAN-OS Release 1.2 or 1.1 after the timer is configured for a VSAN, an error message is issued to warn against strict incompatibilities.

To configure per-VSAN FC timers using Fabric Manager, follow these steps:

Step 1 Choose the VSAN for timer configuration from the Logical Domains pane. If a VSAN is not specified when you change the policies, the changed value is applied to all VSANs in the switch.

Step 2 Expand **Switches**, expand **FC Services** and select **Timers & Policies** in the Physical Attributes tree.

You see timeouts for only switches in the selected VSAN shown in the Information pane.

Step 3 Click **Change Timeouts** to configure the time-out values.

You see the dialog shown in Figure 62-12

Figure 62-12 Change Timeouts per VSAN in Fabric Manager

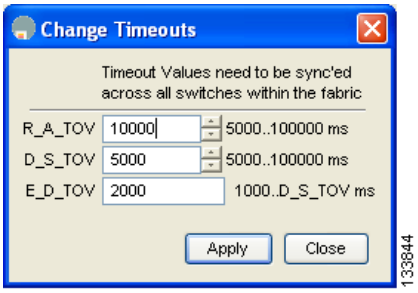

Step 4 Change the timeout values shown inFigure 62-12. Indicate values for R_A_TOV (Resource Allocation Time Out Value), D_S_TOV (Distributed Services Time Out Value), and E_D_TOV (Error Detect Time Out Value).

Step 5 Click **Apply.**

Step 6 Click **Close** to close the dialog box**.**

Configuring a Fabric Analyzer

Fibre Channel protocol analyzers capture, decode, and analyze frames and ordered sets on a link. Existing Fibre Channel analyzers can capture traffic at wire rate speed. They are expensive and support limited frame decoding. Also, to snoop traffic, the existing analyzers disrupt the traffic on the link while the analyzer is inserted into the link.

Cisco has brought protocol analysis within a storage network to a new level with the Cisco Fabric Analyzer. You can capture Fibre Channel control traffic from a switch and decode it without having to disrupt any connectivity, and without having to be local to the point of analysis.

The Cisco Fibre Channel protocol analyzer is based on two popular public-domain software applications:

- libpcap—See http://www.tcpdump.org.
- Ethereal—See http://www.ethereal.com.

Note The Cisco Fabric Analyzer is useful in capturing and decoding control traffic, not data traffic. It is suitable for control path captures, and is not intended for high-speed data path captures.

About the Cisco Fabric Analyzer

The Cisco Fabric Analyzer consists of two separate components (see Figure 62-13):

- Software that runs on the Cisco MDS 9000 Family switch and supports two modes of capture:
  - A text-based analyzer that supports local capture and decodes captured frames
  - A daemon that supports remote capture
- GUI-based client that runs on a host that supports libpcap such as Windows or Linux and communicates with the remote capture daemon in a Cisco MDS 9000 Family switch.

Figure 62-13 Cisco Fabric Analyzer Use

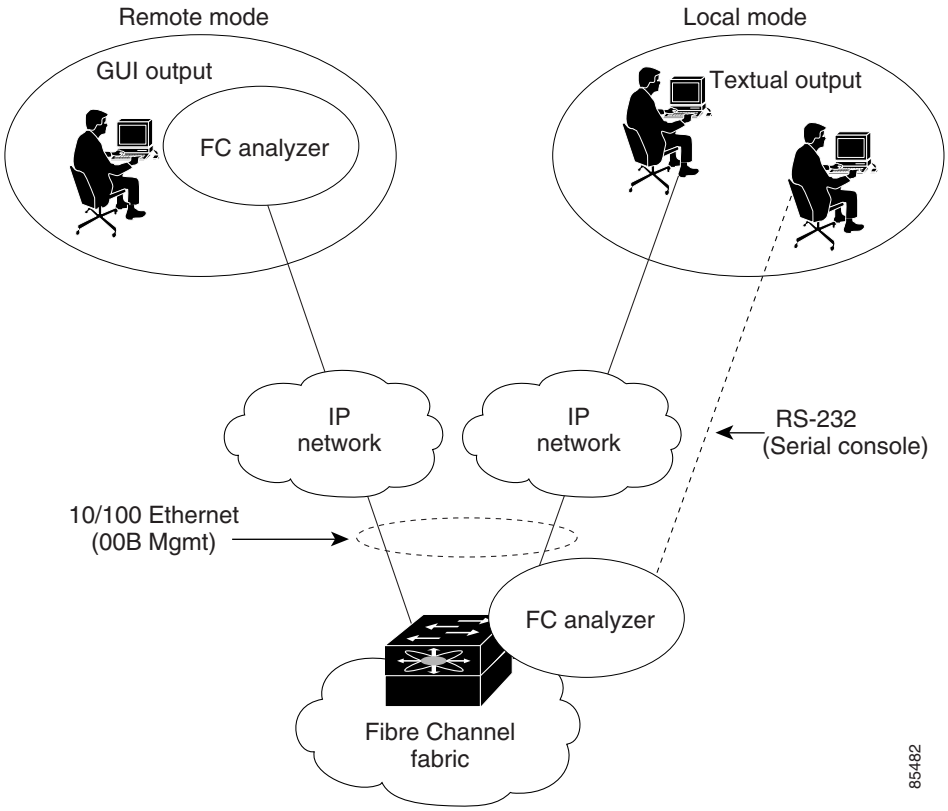

Local Text-Based Capture

This component is a command-line driven text-based interface that captures traffic to and from the supervisor module in a Cisco MDS 9000 Family switch. It is a fully functional decoder that is useful for quick debug purposes or for use when the remote capture daemon is not enabled. Additionally, because this tool is accessed from within the Cisco MDS 9000 Family switch, it is protected by the roles-based policy that limits access in each switch.

Remote Capture Daemon

This daemon is the server end of the remote capture component. The Ethereal analyzer running on a host is the client end. They communicate with each other using the Remote Capture Protocol (RPCAP). RPCAP uses two endpoints, a TCP-based control connection and a TCP or UDP-based data connection based on TCP (default) or UDP. The control connection is used to remotely control the captures (start or stop the capture, or specify capture filters). Remote capture can only be performed to explicitly configured hosts. This technique prevents an unauthorized machine in the network from snooping on the control traffic in the network.

RPCAP supports two setup connection modes based on firewall restrictions.

- Passive mode (default)—The configured host initiates connection to the switch. Multiple hosts can be configured to be in passive mode and multiple hosts can be connected and receive remote captures at the same time.
- Active mode—The switch initiates the connection to a configured host—one host at a time.

Using capture filters, you can limit the amount of traffic that is actually sent to the client. Capture filters are specified at the client end—on Ethereal, not on the switch.

GUI-Based Client

The Ethereal software runs on a host, such as a PC or workstation, and communicates with the remote capture daemon. This software is available in the public domain from http://www.ethereal.com. The Ethereal GUI front-end supports a rich interface such as a colorized display, graphical assists in defining filters, and specific frame searches. These features are documented on Ethereal's website.

While remote capture through Ethereal supports capturing and decoding Fibre Channel frames from a Cisco MDS 9000 Family switch, the host running Ethereal does not require a Fibre Channel connection to the switch. The remote capture daemon running on the switch sends the captured frames over the out-of-band Ethernet management port. This capability allows you to capture and decode Fibre Channel frames from your desktop or laptop.

Configuring the Cisco Fabric Analyzer

You can configure the Cisco Fabric Analyzer to perform one of two captures.

- Local capture—A local capture cannot be saved to persistent storage or synchronized to standby. Launches the textual version on the fabric analyzer directly on the console screen. The capture can also be saved on the local file system.
- Remote capture—A remote capture can be saved to persistent storage. It can be synchronized to the standby supervisor module and a stateless restart can be issued, if required.

To use the Cisco Fabric Analyzer feature, traffic should be flowing to or from the supervisor module.

Sending Captures to Remote IP Addresses

Caution You must use the eth2 interface to capture control traffic on a supervisor module.

To capture remote traffic, use one of the following options:

- The capture interface can be specified in Ethereal as the remote device:

  ```
  rpcap://<ipaddress or switch hostname>/eth2
  ```

  For example:

  ```
  rpcap://cp-16/eth2
  rpcap://17.2.1.1/eth2
  ```

- The capture interface can be specified either in the capture dialog box or by using the -i option at the command line when invoking Ethereal.

  ```
  ethereal -i rpcap://<ipaddress|hostname>[:<port>]/<interface>
  ```

  For example:

  ```
  ethereal -i rpcap://172.22.1.1/eth2
  ```

  or

  ```
  ethereal -i rpcap://customer-switch.customer.com/eth2
  ```

Note For example, in a Windows 2000 setup, click **Start** on your desktop and select **Run**. In the resulting Run window, type the required command line option in the Open field.

Displaying Captured Frames

You can selectively view captured frames by using the display filters feature. For example, instead of viewing all the frames from a capture, you may only want to view Exchange Link Protocol (ELP) request frames. This feature only limits the captured view—it does not affect the captured or the saved frames. Procedures to specify, use, and save display filters are already documented in the Ethereal website (http://www.ethereal.com). Some examples of how you can use this feature are as follows:

- To view all packets in a specified VSAN, use this expression:

  ```
  mdshdr.vsan == 2
  ```

- To view all SW_ILS frames, use this expression:

  ```
  fcswils
  ```

- To view class F frames, use this expression:

  ```
  mdshdr.sof == SOFf
  ```

- To view all FSPF frames, use this expression:

  ```
  swils.opcode == HLO || swils.opcode == LSU || swils.opcode == LSA
  ```

- To view all FLOGI frames, use this expression:

  ```
  fcels.opcode == FLOGI
  ```

- To view all FLOGI frames in VSAN 1, use this expression:

```
fcels.opcode == FLOGI && mdshdr.vsan == 2
```

- To view all name server frames, use this expression:

```
dNS
```

Defining Display Filters

Display filters limit the frames that can be displayed, but not what is captured (similar to any view command). The filters to be displayed can be defined in multiple ways in the GUI application:

- Auto-definition
- Manual definition
- Assisted manual definition
- Only manual definition in local capture
- No assists

Regardless of the definition, each filter must be saved and identified with a name.

Note This GUI-assisted feature is part of Ethereal and you can obtain more information from http://www.ethereal.com.

Capture Filters

You can limit what frames are captured by using the capture filters feature in a remote capture. This feature limits the frames that are captured and sent from the remote switch to the host. For example, you can capture only class F frames. Capture filters are useful in restricting the amount of bandwidth consumed by the remote capture.

Unlike display filters, capture filters restrict a capture to the specified frames. No other frames are visible until you specify a completely new capture.

The syntax for capture filter is different from the syntax for display filters. Capture filters use the Berkeley Packet Filter (BPF) library that is used in conjunction with the libpcap freeware. The list of all valid Fibre Channel capture filter fields are provided later in this section.

Procedures to configure capture filters are already documented in the Ethereal website (http://www.ethereal.com). Some examples of how you can use this feature as follows:

- To capture frames only on a specified VSAN, use this expression:

```
vsan = 1
```

- To capture only class F frames, use this expression:

```
class_f
```

- To capture only class Fibre Channel ELS frames, use this expression:

```
els
```

- To capture only name server frames, use this expression:

  ```
  dns
  ```

- To capture only SCSI command frames, use this expression:

  ```
  fcp_cmd
  ```

Note This feature is part of libpcap and you can obtain more information from http://www.tcpdump.org.

Permitted Capture Filters

This section lists the permitted capture filters.

```
o vsan
o src_port_idx
o dst_port_idx
o sof
o r_ctl
o d_id
o s_id
o type
o seq_id
o seq_cnt
o ox_id
o rx_id
o els
o swils
o fcp_cmd   (FCP Command frames only)
o fcp_data (FCP data frames only)
o fcp_rsp   (FCP response frames only)
o class_f
o bad_fc
o els_cmd
o swils_cmd
o fcp_lun
o fcp_task_mgmt
o fcp_scsi_cmd
o fcp_status
o gs_type      (Generic Services type)
o gs_subtype   (Generic Services subtype)
o gs_cmd
o gs_reason
o gs_reason_expl
o dns   (name server)
o udns (unzoned name server)
o fcs   (fabric configuration server)
o zs    (zone server)
o fc    (use as fc[x:y] where x is offset and y is length to compare)
o els   (use as els[x:y] similar to fc)
o swils (use as swils[x:y] similar to fc)
o fcp   (use as fcp[x:y] similar to fc)
o fcct (use as fcct[x:y] similar to fc)
```

Configuring World Wide Names

The world wide name (WWN) in the switch is equivalent to the Ethernet MAC address. As with the MAC address, you must uniquely associate the WWN to a single device. The principal switch selection and the allocation of domain IDs rely on the WWN. The WWN manager, a process-level manager residing on the switch's supervisor module, assigns WWNs to each switch.

Cisco MDS 9000 Family switches support three network address authority (NAA) address formats (see Table 62-1).

Table 62-1 Standardized NAA WWN Formats

NAA Address	NAA Type	WWN Format	
IEEE 48-bit address	Type 1 = 0001b	000 0000 0000b	48-bit MAC address
IEEE extended	Type 2 = 0010b	Locally assigned	48-bit MAC address
IEEE registered	Type 5 = 0101b	IEEE company ID: 24 bits	VSID: 36 bits

Caution Changes to the world-wide names should be made by an administrator or individual who is completely familiar with switch operations.

Link Initialization WWN Usage

Exchange Link Protocol (ELP) and Exchange Fabric Protocol (EFP) use WWNs during link initialization. The usage details differ based on the Cisco SAN-OS software release:

Both ELPs and EFPs use the VSAN WWN by default during link initialization. However, the ELP usage changes based on the peer switch's usage:

- If the peer switch ELP uses the switch WWN, then the local switch also uses the switch WWN.
- If the peer switch ELP uses the VSAN WWN, then the local switch also uses the VSAN WWN.

Note As of Cisco SAN-OS Release 2.0(2b), the ELP is enhanced to be compliant with FC-SW-3.

Configuring a Secondary MAC Address

To allocate a secondary MAC address, follow these steps:

Step 1 Select a SAN (or a VSAN) from the Logical Domains pane.

You see a list of switches in the Information pane.

Step 2 Expand **Switches**, expand **FC Services** and select **WWN Manager** in the Physical Attributes pane.

Step 3 In the Information pane, scroll until you see the switch on which you want to configure a secondary MAC address (see Figure 62-14).

Figure 62-14 Setting secondary MAC addresses

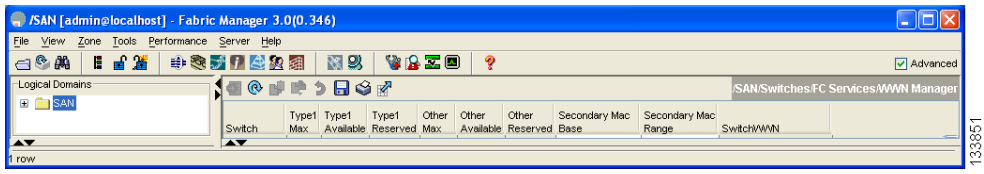

Step 4 Enter the secondary MAC address in the **Secondary Mac Base** field.

Step 5 Enter the range for the secondary MAC address in the **Secondary Mac Range** field.

Step 6 Click the **Apply Changes** icon.

Displaying WWN Information

To display the status of the WWN configuration, follow these steps:

Step 1 Select a SAN (or a VSAN) from the Logical Domains pane.

You see a list of switches in the Information pane.

Step 2 Choose **Switches > FC Services > WWN Manager** from the Physical Attributes pane.

You see the WWN information for each switch in the SAN or VSAN.

FC ID Allocation for HBAs

Fibre Channel standards require a unique FC ID to be allocated to an N port attached to a Fx port in any switch. To conserve the number of FC IDs used, Cisco MDS 9000 Family switches use a special allocation scheme.

Some HBAs do not discover targets that have FC IDs with the same domain and area. Prior to Cisco SAN-OS Release 2.0(1b), the Cisco SAN-OS software maintained a list of tested company IDs which do not exhibit this behavior. These HBAs were allocated with single FC IDs, and for others a full area was allocated.

The FC ID allocation scheme available in Release 1.3 and earlier, allocates a full area to these HBAs. This allocation isolates them to that area and are listed with their pWWN during a fabric login. The allocated FC IDs are cached persistently and are still available in Cisco SAN-OS Release 2.0(1b) (see the "FC ID Allocation for HBAs" section on page 62-23).

As of Cisco SAN-OS Release 2.0(1b), to allow further scalability for switches with numerous ports, the Cisco SAN-OS software maintains a list of HBAs exhibiting this behavior. Each HBA is identified by its company ID (also know as Organizational Unit Identifier, or OUI) used in the pWWN during a fabric log in. Hence a full area is allocated to the N ports with company IDs that are listed and for the others, a single FC ID is allocated. Irrespective of the kind (whole area or single) of FC ID allocated, the FC ID entries remain persistent.

Default Settings

Table 62-2 lists the default settings for the features included in this chapter.

Table 62-2 Default Settings for Advanced Features

Parameters	Default
CIM server	Disabled
CIM server security protocol	HTTP
D_S_TOV	5,000 milliseconds.
E_D_TOV	2,000 milliseconds.
R_A_TOV	10,000 milliseconds.
Time-out period to invoke fctrace	5 seconds.
Number of frame sent by the fcping feature	5 frames.
Remote capture connection protocol	TCP.
Remote capture connection mode	Passive.
Local capture frame limit s	10 frames.
FC ID allocation mode	Auto mode.
Loop monitoring	Disabled.

CHAPTER 63

Management Software FAQ

This chapter answers some of the most frequently asked questions about Cisco Fabric Manager and Device Manager. This chapter contains the following topics:

Installation Issues

When installing Fabric Manager from windows, why does clicking install fail?

To make sure that Java Web Start is installed properly, follow these steps:

Step 1 Go to the Programs menu and see if Java Web Start is there.

Step 2 Start the **Java Web Start** program to make sure there is no problem with the Java Runtime installation.

Step 3 Click the **Preferences** tab, and make sure the proxies settings are fine for Web Start.

Step 4 Check that your browser is set up to handle JNLP settings properly (see the "How do I manually configure a browser for Java Web Start?" section on page 63-6).

If you had older versions of the application and you see an error pop-up window saying cannot open the JNLP file (in the error details), this could be because the Java Web Start cache is messed up. To work around this, clear the cache and retry. To clear the cache, see the "How do I clear the Java Web Start cache?" section on page 63-7.

How do I install Java Web Start on a UNIX machine?

If you are using UNIX (Linux, Solaris) the Sun JRE 1.4.0 and 1.4.1 does not automatically install Java Web Start. However, the Web Start zip file is bundled with the JRE.

To install Java Web Start, follow these steps:

Step 1 Locate the directory where you have installed the JRE. There is a Java Web Start zip file named something like javaws-1_2_linux-i586-i.zip.

Step 2 Unzip this file and run the in**stall.sh** script.

You are prompted to enter the path to the Java installation.

Step 3 Update the mime-type settings for users. This can be done for all users.

```
# /etc/mime.types should contain the line
    type=application/x-java-jnlp-file desc="Java Web Start" exts="jnlp"

# /etc/mailcap should contain the line
    application/x-java-jnlp-file; /javaws %s
```

To install for individual users, add these lines to the files $HOME/.mime.types and $HOME/mailcap.

Why do I have trouble launching Fabric Manager on Solaris?

If you are using Solaris 2.8 and are logged in as root and are using Netscape Navigator 6, you will not be able to register the mime-type. Regular users can register the mime-type with Netscape Navigator 6 by manually adding it. Netscape 4.x works fine for all users.

What do I do if my browser prompts to save JNLP files?

Your browser may not be set up to launch Java Web Start for JNLP mime types. Java Web Start is probably not installed or configured properly (see the "How do I manually configure a browser for Java Web Start?" section on page 63-6).

What do I do if I see a "Java Web Start not detected" error?

If you installed Java Web Start but still see an error message (in red) saying "Java Web Start not detected..." on the switch home page, it could be a simple JavaScript error. We try to detect a Java Web Start installation by running some JavaScript code tested for Internet Explorer and Mozilla (newer versions). On some browsers (for example, Netscape 6.0, Opera) this code does not work properly although the links still work.

- First, try clicking on the install links.
- If that does not work, check to see if the browser helper applications settings are correct (for example, for Netscape 6.0 **Edit > Preferences > Navigator > Helper Applications**). See the "How do I manually configure a browser for Java Web Start?" section on page 63-6.

What do I do if my desktop shortcuts not visible?

For Windows 2000 and Windows NT, we create Program Menu entries (under a new Cisco MDS 9000 program menu) and desktop shortcuts for Fabric Manager and Device Manager. The desktop shortcuts and start menu entries for Fabric Manager and Device Manager are called FabricManager and DeviceManager respectively. In other versions of Windows, including XP, we just create batch files on the desktop called FabricManager.bat and DeviceManager.bat. For UNIX, we create shell scripts called FabricManager.sh and DeviceManager.sh under the $HOME/.cisco_mds9000/bin directory. Note that on Windows, installations run under Mozilla variants of browsers, and the desktop shortcuts do not get created. The workaround is to manually create desktop shortcuts.

How do I upgrade to a newer version of Fabric Manager or Device Manager?

To upgrade to a newer version of Fabric Manager or Device Manager, follow these steps:

Step 1 Close all running instances of Fabric Manager or Device Manager.

Step 2 Point your browser at the switch running the new version and click the appropriate install link. Fabric Manager or Device Manager prompts you to upgrade if the switch is running a newer version.

The installer checks your local copies and updates any newer versions of the software.

How do I downgrade Fabric Manager or Device Manager?

Cisco MDS SAN-OS Release 2.x or later supports downgrades using the installer. For earlier releases, downgrades are not supported through the installer. To downgrade Fabric Manager or Device Manager for an earlier release, you need to manually uninstall them and then install the previous version of Fabric Manager or Device Manager. See the "Downgrading the Management Software" section on page 2-22.

What do I do if an upgrade is not working?

If you are trying to upgrade because Fabric Manager or Device Manager prompted you saying that the switch version is higher, and the upgrade failed, it might be because your default browser settings are incorrect. Some error must have occurred during your last browser upgrade/install. To work around this, launch the browser independently and click on install.

On rare occasions, we have seen the upgrade happen but the version does not change. This is because of HTTP caching in the network. During the upgrade, HTTP requests for files on the switch get cached in the local machine. Even though the switch is in a higher version, the management software installed is at the old version. The workaround for this is to uninstall the Fabric/Device Manager, clear the Java Web Start cache, and then do a clean install.

What do I do if Java Web Start hangs on the download dialog?

To make sure Java Web Start is set up to access the switch in the same way your browser is set up, follow these steps:

Step 1 Start Java Web Start (**javaws.exe** or **javaws**). You see the Java Web Start Application Manager.

Step 2 Choose **File > Preferences > General** and make sure your proxy settings are correct. For example, if you are using an HTTP proxy, set it up here.

Step 3 Choose **Use Browse**r.

Step 4 Click **OK**.

How do I manually configure a browser for Java Web Start?

For browsers like Opera, certain versions of Mozilla, or Konqueror, you must manually register Java Web Start as the helper application for the JNLP files. To do this, the data you need is:

- Description=Java Web Start
- File Extension=jnlp
- Mime Type=application/x-java-jnlp-file
- Application=path-to-javaws (e.g. /usr/local/javaws/javaws)

After setting this up, you may need to restart the browser. If you see "Java Web Start not detected" warnings, you can ignore them. These warnings are based on JavaScript, and not all browsers behave well with JavaScript. Click on the install links to install Fabric Manager or Device Manager.

Note For Windows Users: To set up Java Web Start on *.jnlp files, select **Windows Explorer > Tools > Folder Options > File Types**. Either change the existing setting for JNLP or add one so that *.jnlp files are opened by javaws.exe. This executable is under Program Files\Java Web Start

How do I run Java Web Start from the command line?

If you cannot get your browser to run Java Web Start, you can still run Java Web Start from the command line (javaws.exe or javaws) giving it the URL of the Fabric Manager or Device Manager on the switch as an argument. For example, if your switch IP address is 10.0.0.1, you would use these commands to start Fabric Manager and Device Manager:

```
javaws http://10.0.0.1/cgi-bin/fabric-manager.jnlp
javaws http://10.0.0.1/cgi-bin/element-manager.jnlp
```

What do I do if Windows 2000 crashes (or I see a blue screen)?

Be sure you have Service Pack 3 installed if you are using JRE 1.4.1. (You should actually have been prompted to install Service Pack 3 during the JRE 1.4.1 installation.) If you do not have it installed, Windows 2000 may crash. If you do not want to upgrade to Service Pack 3, you can install JRE 1.4.0.

How do I clear the Java Web Start cache?

To clear the Java Web Start cache, follow these steps:

Step 1 Start the Java Web Start Application Manager (**javaws.exe** or **javaws**).

Step 2 Go to **File > Preferences > Advanced** and clear the applications folder or cache. You can manually delete the .javaws or cache directory. On Windows this is under Documents and Settings, and on UNIX this is under $HOME.

What do I do if my login does not work in Fabric Manager or Device Manager?

Make sure you have done the Initial Setup Routine on the switch. Refer to the *Cisco MDS 9000 Family Configuration Guide*. Quick checks:

- Make sure that the management interface on the switch is up (**show interface mgmt0**).
- Check whether you can connect to the management interface (**ping**).
- Verify the username is valid (**show snmp user**). You can also add/edit the users through the CLI.
- If you have multiple network interfaces, see the "What do I do if Fabric Manager or Device Manager is using the wrong network interface?" section on page 63-9

What do I do if I cannot install Fabric Manager or Device Manager, or run Java, when pcAnyWhere is running?

You can either stop the pcAnyWhere service and install Fabric Manager or Device Manager, or install/update DirectX. For more information,refer to the website at http//forum.java.sun.com/thread.jsp?forum=30&thread=444824&tstart=0&trange=15

What do I do if the Fabric Manager or Performance Manager service shows up as "disabled" in the Services menu?

This could happen if:

- The service menu for Fabric Manager or Performance Manager was open during an uninstall/upgrade.
- The Fabric Manager client or Device Manager was running while doing an uninstall/upgrade.

This error happens when Windows is unable to delete a service completely. A reboot of the host should fix the problem.

What do I do if I am unable to install Fabric Manager or Device Manager, or run Java, when McAfee Internet Suite 6.0 Professional is running?

The McAfee internet suite comes with a virus scanner, firewall, antispam, and privacy management. The privacy management can interfere with the Fabric Manager server-client interactions. To work around this you must shut down the privacy service.

What do I do if I see a ".sm/logon." error displayed when downgrading from MDS SAN-OS Release 2.x (or newer) to 1.3(x)?

The installer does not support a downgrade from Cisco MDS SAN-OS Release 2.x (or newer) to Cisco MDS SAN-OS Release 1.3(x) or earlier. Fabric Manager and Device Manager are backwardly compatible so we suggest running the newer versions. If you still want to downgrade to the lower version, uninstall the 2.0 (or newer) version and then install the older version. For more information, see the "Downgrading the Management Software" section on page 2-22.

General

What do I do if I see errors while monitoring Area chart graphing?

When doing the area chart graphing from the monitor window, if you move the mouse over the Area chart before the first data comes back, you see a java.lang.ArrayIndexOutOfBoundsException error on the message log from JChart getX(). This is because JChart tries to locate a value that does not exist yet. This might be fixed in a future version of JChart.

What do I do if I see "gen error" messages?

Usually a "gen error" means that the SNMP agent on the switch had an unexpected error in the process of serving an SNMP request. However, when you are accessing the switch through a VPN connection or any sort of NAT scheme, all errors are reported as gen error. This is a known problem and will be fixed in a future release. You can verify whether this was the reason behind your gen error by trying to reproduce this error in an environment where there is no network address translation (where you are on the same network as the switch).

What do I do if disk images in the Device Manager Summary View are not visible?

On some occasions the Summary View table in the Device Manager does not show the icons for disks attached to a Fx port. This is because the FC4 features are empty for this port. A LUN discovery must be issued to discover information about these hosts/disks that do not register their FC4 types. You can do this in the Device Manager by clicking **FC > Advanced > LUNs**.

What do I do if I am unable to set both the D_S_TOV and E_D_TOV timers in Device Manager?

If you modify both E_D_TOV and D_S_TOV at the same time, and the new D_S_TOV value is larger than the old E_D_TOV value, you will get a WrongValue error. To work around this, you must change the values separately.

What do I do if columns in Device Manager tables are too small?

If Device Manager is trying to display a large table and your switch is running slowly, the table will come up with the tabs being hidden. To work around this, you must resize the window to see the data.

What do I do if fabric changes are not propagated onto the map (for example, links don't disappear)?

Fabric Manager shows that a device or port is down by displaying a red cross on that port or device. However, Fabric Manager does not remove any information that's already discovered. You must rediscover to correctly update the map.

What do I do if the PortChannel creation dialog becomes too small after several uses?

After several uses, the MemberList TextBox (in the PortChannel Create Window) does not display as it should. It changes from a long TextBox with a ComboBox for choosing ports, to a small square TextBox that is too small to choose ports. This is a known problem and will be fixed in a future release. To work around this problem, stop and restart Fabric Manager or Device Manager.

What do I do if I see errors after IPFC configuration?

When IPFC and out of band management are configured, the Device Manager might not work using SNMPv3 if you use the IPFC address. The workaround is either to use the management interface (mgmt0) address, or to use SNMPv1/v2c over IPFC.

What do I do if Fabric Manager or Device Manager is using the wrong network interface?

The problem happens because the underlying Java library picks a local interface arbitrarily. To work around this, supply a command line argument before starting the Fabric/Device Manager. In the desktop shortcut or shell script or batch file, add the following parameter "-Device Managerds.nmsAddress="

For example, in Windows the line looks like ".javaw.exe -Device Managerds.nmsAddress=X.X.X.X -cp .".

In desktop shortcuts, this length could exceed the maximum characters allowed. If this happens, delete the "-Dsun.java2d.ddoffscreen=false" portion to make more space. Newer versions of Fabric Manager (Release 1.2 and later) allow you to pick a preferred network interface.

What do I do if I see display anomalies in Fabric Manager or Device Manager?

If you see Fabric Manager or Device Manager submenus detached from menus, the mouse pointer in Fabric Manager Map is slow to react to mouse movement, or a wrong tooltip is displayed, these are display anomalies, not problems with Fabric Manager or Device Manager.

Some older video cards exhibit these display anomalies. To fix this, first try updating the video drivers. If this doesn't solve the problem, replace the video card.

What do I do if most of my Physical Attributes catagories disappear?

You have somehow turned off advanced features. Look for the check box Advanced Features in the upper right of the Fabric Manager screen. Check the box.

What do I do if I can't see the Information pane?

The information pane should be in the upper half of the screen above the map in Fabric Manager. The map my be covering it. Drag the edge of the map window down or use the black triangles to reorganize the display.

Why is the active zone set in edit zone always shown in bold (even after successful activation)?

A member of this VSAN must be participating in IVR zoning. Because the IVR zones get added to active zones, the active zone set configuration is always different from the local zone set configuration with the same name. The zone set name is always bold.

Can I create a zone with prefix IVRZ or a zone set with name nozonset?

Do not use these special names. These names are used by the system for identifying IVR zones.

What do I do when One-Click License Install fails, and I cannot connect to the Cisco website?

The one-click license install tries to open an HTTP connection to the Cisco website. If you do your browsing using an HTTP proxy then the following command- line variables need to be added to your Fabric Manager client scripts:

```
-Dhttps.proxyHost and -Dhttps.proxyPort.
```

In case your one-click install URL starts with "http://" (and not "https://"), the variables are:

```
-Dhttp.proxyHost and -Dhttp.proxyPort.
```

For example, in Windows, edit the MDS 9000\bin\FabricManager.bat file and add to the JVMARGS "-Dhttps.proxyHost=HOSTADDRESS -Dhttps.proxyPort=HOSTPORT".

What do I do when Fabric Manager client and Device Manager cannot connect to the switch?

Fabric Manager or Device Manager using SNMPv3 at Cisco MDS SAN-OS Release 1.3(3) or earlier can't manage a switch running Release 1.3(4) or later. This might affect a software upgrade using Fabric Manager from Release 1.3(3) to Release 1.3(4).

What do I do when the License Wizard fails to fetch license keys, saying connect failed?

Java versions 1.4.2_01 and older don't seem to have the right set of CA (certifying authority) certificates to validate the SSL certificates on the EMC server (https).The license wizard is unable to make an https connection to the EMC servers. The workaround is to install the latest 1.4(x) version of Java, preferably 1.4.2_04 or later.

How do I increase the log window size in Fabric Manager Client?

To limit the memory usage by FM Client, the log window is limited to 500 lines by default. If you want to increase this, edit sm.properties in < *install directory*>/db/<*user*> directory and change LogBufferSize.

When do I do when the FM Server Database fails to start or has a file locking error?

In the database log (FMPersist.log) you will see an error message "The database is already in use by another process". The HsqlDB 1.7.1 version has this problem. The file lock problem seems to happen occasionally, and can be resolved by shutdown and restart of the db server. On windows this can be done by stopping and starting the FMPersist service and on Unix just run the FMPersist.sh script with the argument restart.

Windows Issues

What do I do when text fields show up too small, and I cannot enter any data?

When Reflection X is running, certain text fields in the Fabric Manager and Device Manager are not rendered to the full width of the field. Resize the dialog box to see the text fields properly.

What do I do when CiscoWorks fails to start in the browser?

CiscoWorks fails to come up in the browser. This could be because CiscoWorks does not support Java JVM 1.4.0. To turn off Java JVM 1.4.0 in Internet Explorer, select the **Tools/Internet Options** menu item, click on the **Advanced** tab, and uncheck the **Use Java 2 1.4.0** option.

What do I do when printing causes an application crash?

On Windows NT there is a known Sun JVM bug - the printservice crashes the VM. The solution suggested by Sun is to update NT with SP 6. For more details refer to:
http://developer.java.sun.com/developer/bugParade/bugs/4530428.html.

What do I do when Windows XP hangs (or I see a blue screen)?

Windows XP with the ATI Radeon AGP graphics cards has known to freeze (hang) when a Java application exits. The newer drivers from ATI seem to have fixed this problem. The other workaround is to run the application with "-Dsun.java2d.noddraw=true". We do this today in the shortcut and shell scripts we create. For more details refer to:
http://developer.java.sun.com/developer/bugParade/bugs/4713003.html.

What do I do when Fabric Manager and Device Manager icons disappear?

On certain versions of Windows, certain images disappear. This is a Java bug. We have a workaround that is already in place (disable DirectDraw acceleration) - but there are still cases where this problem might arise. For more details refer to:
http://developer.java.sun.com/developer/bugParade/bugs/4664818.html.

What do I do when Fabric Manager hangs when dragging an existing zone member to a zone?

When dragging a zone member to a zone (where that member is already present) you get an error message saying the zone member is already present and the application freezes. This is a Sun Java bug, and the problem is seen with JRE versions earlier than 1.4.2. For more details refer to http://developer.java.sun.com/developer/bugParade/bugs/4633417.html. Use a Sun JRE with version 1.4.2 or later where this problem does not occur.

What do I do when Device Manager or Fabric Manager window content disappears in Windows XP?

Device Manager or Fabric Manager main window content disappears in Windows XP due to a Java bug. Refer to the following website:
http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=4919780.

Minimize or maximize the window and restore to the normal size to restore the window content. Disabling Direct Draw may also prevent this from happening by adding "-Dsun.java2d.noddraw=true" to JVMARGS in <*FM-install-dir*>/bin/FabricManager.bat and DeviceManager.bat

What do I do when SCP/SFTP fails when a file is copied from local machine to the switch?

If there are embedded spaces in the file path, then windows scp/sftp might fail. You will get a copyDeviceBusy error from the switch. In tools such as the License Wizard either make sure tftp copy can be done or pick filenames with no spaces.

UNIX Issues

What do I do when the parent menus disappear?

Displaying a submenu may occasionally cause the parent menu to disappear. For more details on this bug, refer to: http://developer.java.sun.com/developer/bugParade/bugs/4470374.html.

What do I do when the web browser cannot find web server even it is running?

This can happens when web browser uses proxy server. To check that for Internet Explorer, choose tools in menu, then choose internet options, then choose connection subpanel, then click Lan Setting. A dialog comes up, verify the proxy setting.

How do I fix a "too many open files" error?

If you are running the JVM (Java Virtual Machine) on Linux and the drive where Java is installed or your home directory is NFS mounted, there is an open bug against the Sun JDK about errors acquiring file locks. The symptoms for the Fabric Manager are that launching a Device Manager or saving/opening files will fail, giving a "too many open files" I/O or socket exception. The JVM keeps trying to open a file on the NFS mounted drives, fails, and keeps trying to do it until it hits the 1024 file descriptors limit. Workarounds (assuming /tmp is a local disk - replace it with your tmp area):

- System Preferences

  Make sure the system level preferences are stored on a local disk. The system preferences are stored in $JAVA_HOME/.systemPrefs where JAVA_HOME is where you have installed the JDK. If this directory is NFS mounted, then just do the following:

  ```
  $ rm -rf $JAVA_HOME/.systemPrefs<
  $ mkdir /tmp/.systemPrefs
  $ ln -s /tmp/.systemPrefs $JAVA_HOME/.systemPrefs
  ```

  The problem with this workaround is that you have to make sure /tmp/.systemPrefs exists on every box where you are using $JAVA_HOME. We recommend installing the JVM as root and on a local disk.

- User Preferences

  If your home directory is NFS mounted and you are getting this problem. Do the following:

  ```
  $ rm -rf $HOME/.java
  $ mkdir /tmp/.java.$USER
  $ ln -s /tmp/.java.$USER $HOME/.java
  ```

For further details, see the following URLs:
http://developer.java.sun.com/developer/bugParade/bugs/4673298.html
http://developer.java.sun.com/developer/bugParade/bugs/4635353.html

Other

How do I set the map layout so it stays after Fabric Manager restarted?

If you have arranged the map to your liking and would like to "freeze" the map so that the objects stay as they are even after you stop Fabric Manager and restart it again, follow these steps:

Step 1 Right-click in a blank space in the map. You see a menu.

Step 2 Select **Layout > Fix All Nodes** from the menu.

What do I do when two switches show on the map, but there is only one switch?

If two switches show on your map, but you only have one switch, it may be that you have two switches in a non-contiguous VSAN that have the same Domain ID. Fabric Manager uses <vsanId><domainId> to look up a switch, and this can cause the fabric discovery to assign links incorrectly between these errant switches.

The workaround is to verify that all switches use unique domain IDs within the same VSAN in a physically connected fabric. (The fabric configuration checker will do this task.)

What does a red/orange/dotted line through the switch mean?

If a red line shows through your switch, this means Fabric Manger sees something wrong with the switch. Choose **Switches** in the Physical Attributes pane to see a status report in the information pane. A module, fan, or power supply has failed or is offline and plugged in.

If a dotted orange line shows through your switch, this indicates a minor status warning for that switch. Usually it means an issue with one of the modules. The tooltip should say exactly what is wrong. Hold the mouse over the switch to see the tooltip.

Below are tables of color settings and tooltip definitions for Fabric Manager and Device Manager.

Table 63-1 Fabric Manager and Device Manager Color Definitions

Fabric Manager Color	Definition
Red Slash	Cannot communicate with a switch via SNMP.
Red X	Cannot communicate with or see a switch in the Domain Manager/Fabric Configuration Server list of fabric switches.
Device Manager Color	**Definition**

Table 63-1 Fabric Manager and Device Manager Color Definitions (continued)

Fabric Manager Color	Definition
Green Square with Mode (e.g., F, T, TE, U/I for FICON)	Port up.
Orange Square with Mode	Trunk incomplete.
Orange Cross	Ols or Nos received.
Brown Square	Port is administratively down.
Light Gray Square	Port is not manageable.
Red Cross	HardwareFailure/LoopbackDiagFailure/LinkFailure
Red Square	Any other kind of configuration failure.
No Square or Black Square	Port not yet configured.

Table 63-2 Device Manager Tooltip Definitions

Tooltip	Definition
adminDown	The port is administratively down.
bitErrRTThresExceeded	Bit error rate too high.
bundleMisCfg	Misconfiguration in PortChannel membership detected.
channelAdminDown	This port is a member of a PortChannel and that PortChannel is administratively down.
channelConfigurationInProgress	This port is undergoing a PortChannel configuration.
channelOperSuspended	This port is a member of a PortChannel and its operational parameters are incompatible with the PortChannel parameters.
deniedDueToPortBinding	Suspended due to port binding.
domainAddrAssignFailureIsolation	The elected principal switch is not capable of performing domain address manager functions so no Nx_port traffic can be forwarded across switches, hence all Interconnect_Ports in the switch are isolated.
domainInvalidRCFReceived	Invalid RCF received.
domainManagerDisabled	Domain manager is disabled.
domainMaxReTxFailure	Domain manager failure after maximum retries.
domainOtherSideEportIsolation	The peer E port is isolated.
domainOverlapIsolation	There is a overlap in domains while attempting to connect two existing fabrics.
elpFailureClassFParamErr	Isolated for ELP failure due to class F parameter error.

Table 63-2 Device Manager Tooltip Definitions (continued)

Tooltip	Definition
elpFailureClassNParamErr	Isolated for ELP failure due to class N parameter error.
elpFailureInvalidFlowCTLParam	Isolated for ELP failure due to invalid flow control parameter.
elpFailureInvalidPayloadSize	Isolated for ELP failure due to invalid payload size.
elpFailureInvalidPortName	Isolated for ELP failure due to invalid port name.
elpFailureInvalidSwitchName	Isolated for ELP failure due to invalid switch name.
elpFailureInvalidTxBBCredit	Isolated for ELP failure due to invalid transmit B2B credit.
elpFailureIsolation	During a port initialization the prospective Interconnect_Ports find incompatible link parameters.
elpFailureLoopbackDetected	Isolated for ELP failure due to loopback detected.
elpFailureRatovEdtovMismatch	Isolated for ELP failure due to R_A_TOV or E_D_TOV mismatch.
elpFailureRevMismatch	Isolated for ELP failure due to revision mismatch.
elpFailureUnknownFlowCTLCode	Isolated for ELP failure due to invalid flow control code.
ePortProhibited	Port down because FICON prohibit mask in place for E/TE port.
eppFailure	Trunk negotiation protocol failure after maximum retries.
errorDisabled	The port is not operational due to some error conditions that require administrative attention.
escFailureIsolation	During a port initialization the prospective Interconnect_Ports are unable to proceed with initialization as a result of Exchange Switch Capabilities (ESC).
fabricBindingDBMismatch	fabric bindingactive database mismatch with peer.
fabricBindingDomainInvalid	Peer domain ID is invalid in fabric binding active database.
fabricBindingNoRspFromPeer	Fabric binding no response from peer.
fabricBindingSWWNNotFound	Peer switch WWN not found in fabric binding active database.
fcipPortAdminCfgChange	FCIP port went down due to configuration change.
fcipPortKeepAliveTimerExpire	FCIP port went down due to TCP keep alive timer expired.

Table 63-2 Device Manager Tooltip Definitions (continued)

Tooltip	Definition
fcipPortMaxReTx	FCIP port went down due to max TCP retransmissions reached the configured limit.
fcipPortPersistTimerExpire	FCIP port went down due to TCP persist timer expired.
fcipPortSrcAdminDown	FCIP port went down because the source ethernet link was administratively shutdown.
fcipPortSrcLinkDown	FCIP port went down due to ethernet link down.
fcipSrcModuleNotOnline	FCIP port went down due to source module not online.
fcipSrcPortRemoved	FCIP port went down due to source port removal.
fcotChksumErr	FSP SPROM checksum error.
fcotNotPresent	SFP (GBIC) not present.
fcotVendorNotSupported	FSP (GBIC) vendor is not supported.
fcspAuthenfailure	Fibre Channel security protocol authorization failed.
ficonBeingEnabled	FICON is being enabled.
ficonNoPortnumber	No FICON port number.
ficonNotEnabled	FICON not enabled.
ficonVsanDown	FICON VSAN is down.
firstPortNotUp	In a over subscribed line card, first port cannot be brought up in E mode when the other ports in the group are up.
firstPortUpAsEport	In a over subscribed line card, when the first port in a group is up in E mode, other ports in that group cannot be brought up.
hwFailure	Hardware failure.
incomAdminRxBBCreditPerBuf	Disabled due to incompatible admin port rxbbcredit, performance buffers.
incompatibleAdminMode	Port admin mode is incompatible with port capabilities.
incompatibleAdminRxBBCredit	Receive BB credit is incompatible.
incompatibleAdminRxBufferSize	Receive buffer size is incompatible.
incompatibleadminSpeed	Port speed is incompatible with port capabilities.
initializing	The port is being initialized.
interfaceRemoved	Interface is being removed.
invalidAttachment	Invalid attachment.
invalidConfig	This port has a misconfiguration with respect to port channels.
invalidFabricBindExh	Invalid fabric binding exchange.

Table 63-2 Device Manager Tooltip Definitions (continued)

Tooltip	Definition
linkFailCreditLoss	Link failure due to excessive credit loss indications.
linkFailCreditLossB2B	Link failure when link reset (LR) operation fails due to queue not empty.
linkFailDebounceTimeout	Link failure due to re-negotiation failed.
linkFailLineCardPortShutdown	Link failure due to port shutdown.
linkFailLinkReset	Link failure due to link reset.
linkFailLIPF8Rcvd	Link failure due to F8 LIP received.
linkFailLIPRcvdB2B	Link failure when loop initialization (LIP) operation fails due to non empty receive queue.
linkFailLossOfSignal	Link failure due to loss of signal.
linkFailLossOfSync	Link failure due to loss of sync.
linkFailLRRcvdB2B	Link failure when link reset (LR) operation fails due to non-empty receive queue.
linkFailNOSRcvd	Link failure due to non-operational sequences received.
linkFailOLSRcvd	Link failure due to offline sequences received.
linkFailOPNyRETB2B	Link failure due to open primitive signal returned while receive queue not empty.
linkFailOPNyTMOB2B	Link failure due to open primitive signal timeout while receive queue not empty.
linkFailPortInitFail	Link failure due to port initialization failure.
linkFailPortUnusable	Link failure due to port unusable.
linkFailRxQOverFlow	Link failure due to receive queue overflow.
linkFailTooManyINTR	Link failure due to excessive port interrupts.
linkFailure	Physical link failure.
loopbackDiagFailure	Loopback diagnostics failure.
loopbackIsolation	Port is connected to another port in the same switch.
noCommonVsanIsolation	Trunk is isolated because there are no common vsans with peer.
none	No failure.
nonParticipating	During loop initialization, the port is not allowed to participate in loop operations
offline	Physical link is in offline state as defined in the FC-FS standards.
ohmsExtLBTest	Link suspended due to external loopback diagnostics failure.
other	Undefined reason.

Table 63-2 Device Manager Tooltip Definitions (continued)

Tooltip	Definition
parentDown	The physical port to which this interface is bound is down.
peerFCIPPortClosedConnection	Port went down because peer FCIP port closed TCP connection.
peerFCIPPortResetConnection	Port went down because the TCP connection was reset by the peer FCIP port.
portBindFailure	Port got isolated due to port bind failure.
portBlocked	Port blocked due to FICON.
portChannelMembersDown	No operational members.
portFabricBindFailure	Port isolated due to fabric bind failure.
portGracefulShutdown	Port shutdown gracefully.
portVsanMismatchIsolation	An attempt is made to connect two switches using non-trunking ports having different port VSANs.
rcfInProgres	An isolated xE_port is transmitting a reconfigure fabric, requesting a disruptive reconfiguration in an attempt to build a single, non-isolated fabric. Only the Interconnect_Ports can become isolated.
srcPortNotBound	No source port is specified for this interface.
suspendedByMode	Port that belongs to a port channel is suspended due to incompatible operational mode.
suspendedBySpeed	Port that belongs to a port channel is suspended due to incompatible operational speed.
suspendedByWWN	Port that belongs to a port channel is suspended due to incompatible remote switch WWN.
swFailure	Software failure.
tooManyInvalidFLOGIs	Suspended due to too many invalid FLOGIs.
tovMismatch	Link isolation due to TOV mismatch
trunkNotFullyActive	Some of the VSANs which are common with the peer are not up.
upgradeInProgress	Line card upgrade in progress.
vsanInactive	Port VSAN is inactive. The port becomes operational again when the port VSAN is active.
vsanMismatchIsolation	This VSAN is not configured on both sides of a trunk port.
zoneMergeFailureIsolation	The two Interconnect_Ports cannot merge zoning configuration after having exchanged merging request for zoning.
zoneRemoteNoRespIsolation	Isolation due to remote zone server not responding.

How do I upgrade without losing map settings?

When you upgrade from one version of Fabric Manager to another, there is a way to prevent the loss of map settings (enclosure names, placement on the map, etc.)

The MDS 9000/db directory contains subfolders for each user (and one for fmserver). In these subfolders are files for all discovered fabrics (*.dat) and maps (*.map). These are upgradable between versions. If you need to clear the fabric cache, you should first export the enclosures to a file to avoid losing them. Everything else aside from enclosures and map coordinates are stored on the switch. The preferences, last opened, and site_ouis.txt format doesn't change from release to release.

How do I preserve historical data when moving Fabric Manager server to new host?

To preserve your data when moving Fabric Manager Server to a new host, follow these steps:

Step 1 Copy the cisco_mds9500/pm directory from the old host to the new host. Place it in the MDS 9000 directory (on a Windows PC, the default installation location for this directory is C:\Program Files\Cisco Systems\MDS 9000).

Step 2 On the new host, run **PMUpgrade.bat** from the MDS 9000\bin folder. This creates files and a new directory structure. There is a directory for each switch for which you have collected data.

Step 3 Continue to collect data on a specific switch by copying the **db** subfolder from that switch's folder to the **pm** folder.

Step 4 On the new host, restart the Performance Manager Service (Windows) or Daemon (UNIX). You can use the **bin/PM.bat** file to do this, or you can choose **Performance > Collector > Restart** from the Fabric Manager menu.

Step 5 Export the enclosures to a file.

Step 6 Re-import the enclosures on the new host.

Step 7 Be sure to turn off the original service on the old host.

Are there restrictions when using Fabric Manager across FCIP?

Fabric Manager will work with no restriction across an FCIP tunnel, as long as the tunnel is up. However, Fabric Manager cannot automatically discover a Cisco SN5428 mgmt IP address in the fabric. For that switch, it will display a red slash through an FCIP device because of a timeout error. It will still see all targets, initiators, and ISLs attached to a Cisco SN5428 (or any other switch) as long as they appear in the name server or FSPF.

To work around this, you can manually enter the IP address in the Switches table, and click Apply. If the community string is correct, the red slash will go away. Even if the community string is incorrect, double-clicking on the Cisco SN5428 will launch the web tool.

How do I fix a "Please insure that FM server is running on localhost" message?

You may see this error message if you cannot connect to the fabric and your PC has multiple network interface cards. The problem may be that Fabric Manager is trying to communicate through the wrong interface (you can verify this by checking the FMServer.log file).

Generally it is best to let Fabric Manager choose the interface on startup. If you are getting the above error, something may have gone wrong.

To reset Fabric Manager so that it chooses the interface next time it starts, follow these steps:

Step 1 Open the server.properties file in the Fabric Manager installation directory. On a Windows platform, this file is in C:\Program Files\Cisco Systems\MDS 9000 by default.

Step 2 Comment out the line: snmp.localaddress.

Step 3 Save and exit the file.

Step 4 Restart Fabric Manager.

Note There are some cases where you would not want to do this, and should manually select the interface that Fabric Manager uses. For more information, see the"How do I run Cisco Fabric Manager with multiple interfaces?" section on page 63-21.

How do I run Cisco Fabric Manager with multiple interfaces?

If your PC has multiple interfaces (NICs), the four Cisco Fabric Manager applications detect these interfaces automatically (ignoring loopback interfaces). Fabric Manager Client and Device Manager detect all interfaces on your PC each time you launch them, and allow you to select one. Fabric Manager Server and Performance Manager detect on initial install, and allows you to select one. You are not prompted again to choose an interface with these two applications.

There may be circumstances where you will want to change the interface you are using. For example:

- If you add an interface after you have installed Fabric Manager Server and/or Performance Manager
- If you decide to use a different interface than the one you initially selected
- If for any reason one of the Cisco Fabric Manager applications did not detect multiple interfaces

Refer to the following sections, depending on which application you want to recognize the interface.

- Manually specifying an interface for Fabric Manager Server, page 63-21
- Manually specifying an interface for Fabric Manager Client or Device Manager, page 63-22

Manually specifying an interface for Fabric Manager Server

To specify an interface for Fabric Manager Server (including Performance Manager and Fabric Manager Web Services), follow these steps:

Step 1 Go to the MDS 9000 folder. On a Windows platform, this folder is at C:\Program Files\Cisco Systems\MDS 9000 by default.

Step 2 Edit the server.properties file with a text editor.

Step 3 Scroll until you find the line: snmp.localaddress.

Step 4 If the line is commented, remove the comment character.

Step 5 Set the local address value to the IP address or interface name of the NIC you want to use.

Step 6 Save the file.

Step 7 Stop and restart Fabric Manager Server.

Manually specifying an interface for Fabric Manager Client or Device Manager

To specify an interface for the Fabric Manager Client or Device Manager, follow these steps:

Step 1 Go to the MDS 9000/bin folder. On a Windows platform, this folder is at C:\Program Files\Cisco Systems\MDS 9000 by default.

Step 2 Edit the DeviceManager.bat file or the FabricManager.bat file.

Step 3 Scroll to the line that begins with set JVMARGS=

Step 4 Add the parameter -Device Managerds.nmsaddress=*ADDRESS*, where *ADDRESS* is the IP address or interface name of the NIC you want to use.

Step 5 Save the file and relaunch Fabric Manager Client or Device Manager.

How do I configure an HTTP proxy server?

If your network uses a proxy server for HTTP requests, make sure the Java Web Start Application Manager is properly configured with the IP address of your proxy server.

To configure a proxy server in the Java Web Start Application Manager, follow these steps:

Step 1 Launch the Java Web Start application.

Step 2 Choose **File > Preferences** from the Java WebStart Application Manager.

Step 3 Choose the **Manual** radio button and enter the IP address of the proxy server in the HTTP Proxy field.

Step 4 Enter the HTTP port number used by your proxy service in the **HTTP Port** field.

Step 5 Click **OK**.

How do I clear the topology map?

If you have a switch that you have removed from the fabric, there will be a red X through the switch's icon. You can clear this information from the Fabric Manager client, or from the Fabric Manager server (which will clear the information for all clients) without having to reboot the switch.

To clear information from topology maps using Fabric Manager, follow these steps:

Step 1 Click the **Refresh Map** icon in the Fabric pane.

This clears the information from the client.

Step 2 Click **Purge Down Elements** in the Server menu.

This clears the information from the server.

Caution Any devices not currently accessible (may be offline) are purged.

How can I use Fabric Manager in a mixed software environment?

You can use Fabric Manager version 2.0(x) to manage a mixed fabric of Cisco MDS 9000 switches. Certain 2.0 feature tabs will be empty for any switches running a software version that does not support those features.

How do I fix a "corrupted jar file" error when launching Fabric Manager?

If you get the following error:

```
An error occurred while launching the application Fabric Manager.

download error:corrupted jar file at <ipaddress>\Device Managerboot.jar
```

(Where <*ipaddress*> is that of the switch)

The error message you are getting indicates that the Java Web Start cache is corrupted. You can try clearing your Java Web Start cache first. To clear the Cache either run Java Web Start (from the Programs menu) and under the **preferences** select **clear cache**. Or do it manually by first making sure all Fabric Manager or Device Manager instances are closed and then deleting .javaws/cache. In the newer JREs this directory is created under Documents and Settings\USERNAME and in the older ones it used to be under Program Files\Java Web Start.

You can also browse beneath the cache folder and delete the offending IPAddress folder (e.g. cache/http/D10.0.0.1).

Also, check to make sure that the host is not running a virus checker / java blocker?

Also you can run the un-install program, then deleting .cisco_mds directory. Then re-install Fabric Manager.

How do I search for devices in a fabric?

In Fabric Manager, it is possible to search for one or more devices by different attributes, including pWWN.

To perform a search in Fabric Manager, follow these steps:

Step 1 Right-click the map and choose **Find Elements** from the drop-down menu.

You see the Find Fabric dialog box.

Step 2 Choose **End Device** from the first drop-down list (see Figure 63-1).

Step 3 Choose **Port WWN** from the first drop-down list (see Figure 63-1).

You can also enter only part of the WWN and use a wildcard (*) character (for example, you can enter ***fb*f8**).

Figure 63-1 Find Fabric Dialog box with End Device and Port WWN selected

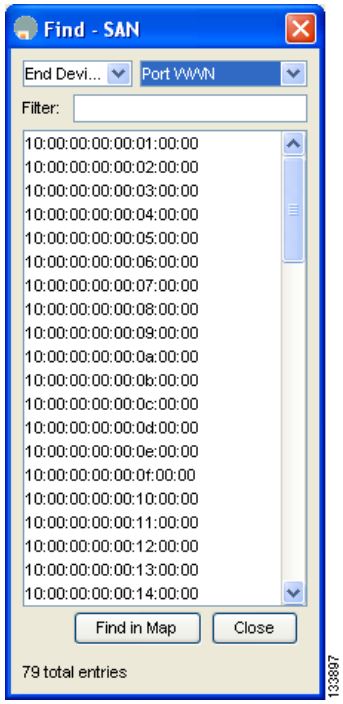

Step 4 Click **Find in Map** (see Figure 63-1).

You see the device highlighted in the Fabric pane. Right-click any device to see the attributes for that device. You can also select a link leading to a device to see the attributes for the link.

How does Fabric Manager Server licensing work?

You must install a Cisco MDS 9000 Family Cisco Fabric Manager Server package on at least one switch in each fabric where you intend to manage switches, if you intend to use the enhanced management capabilities the license package provides. You must also license all switches you plan to monitor with

the Performance Manager (historical performance monitoring) feature. Failure to license all switches can prevent effective use of the Flow performance monitoring, so it is recommended to license all switches in each fabric managed by Cisco Fabric Manager Server.

You are free to try Cisco Fabric Manager Server capabilities prior to installing a license, but the those extended functions will stop working after the 120-day grace period expires. Standard Cisco Fabric Manager configuration and management capabilities will continue to be accessible without any licensed switches after the grace period expires.

How do I manage multiple fabrics?

To monitor and manage multiple fabrics, you must persist one or more fabrics. Do this by checking the **Persist** checkbox on the **Server>Admin** dialog Fabric tab. You must also use switches running SAN-OS Release 1.3.x or greater in both fabrics, and you must use the same user/password on both fabrics. Both fabrics must not be physically connected.

How can I clear an orange X through a switch caused by license expiration?

If you are using a licensed feature and that license is allowed to expire, Fabric Manager shows a license violation, and an orange X is placed through the switch on the Fabric Manager map.

To clear the license violation message and the orange X, stop the Cisco Fabric Manager service on the host, and restart it again.

CHAPTER 64

Monitoring System Processes and Logs

This chapter provides details on monitoring the health of the switch. It includes the following sections:

Displaying System Processes

To obtain general information about all processes using Device Manager, follow these steps:

Step 1 Select **Running Processes** from the Admin menu.

You see the screen in Figure 64-1.

Figure 64-1 Running Processes

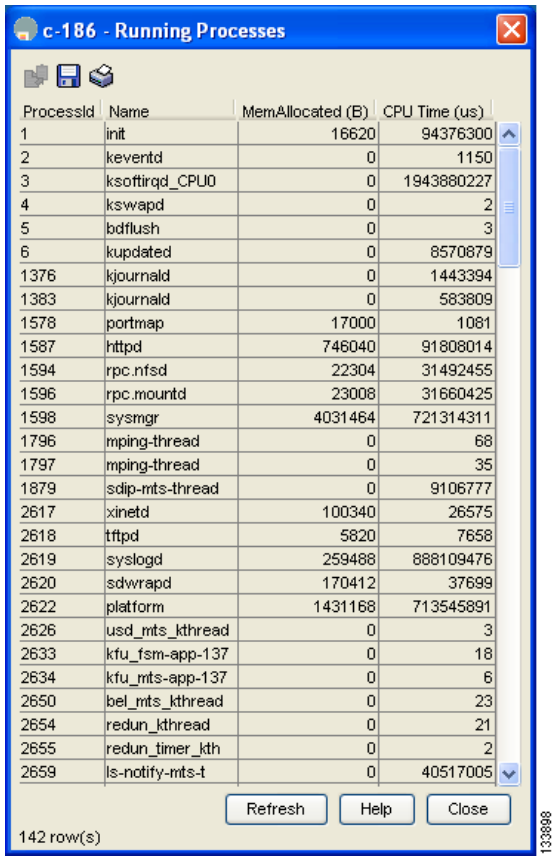

Where:

- ProcessId = Process ID.
- Name = Name of the process
- MemAllocated = Sum of all the dynamically allocated memory that this process has received from the system, including memory that may have been returned
- Runtime (ms) = CPU time the process has used, in microseconds

Displaying System Status

To display system status from Device Manager, follow these steps:

Step 1 Select **System** from the Physical menu.

You see the dialog box in Figure 64-2.

Figure 64-2 System Dialog Box

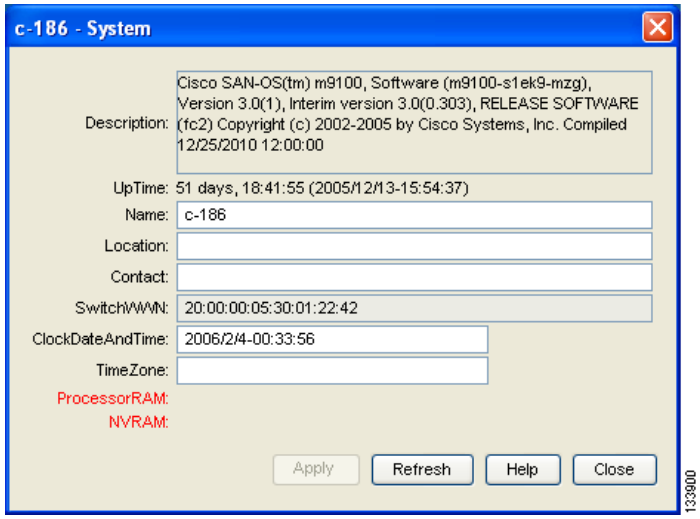

Core and Log Files

For information on copying core and log files, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

Displaying Core Status

To display cores on a switch using Device Manager, follow these steps:

Step 1 Be sure SSH2 is enabled on this switch.

Step 2 Select **Show Cores** from the Admin menu.

You see the dialog box in Figure 64-3.

Figure 64-3 Show Cores Dialog

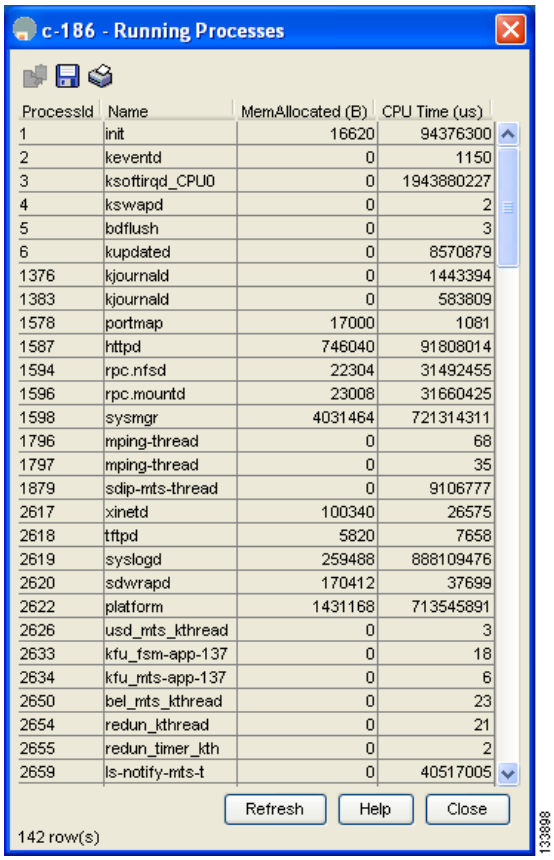

Where:

`Module-num` shows the slot number on which the core was generated.In this example, the `fspf core` was generated on the active supervisor module (slot 5), `fcc` was generated on the standby supervisor module (slot 6), and `acltcam` and `fib` were generated on the switching module (slot 8).

Clearing the Core Directory

To display cores on a switch using Device Manager, follow these steps:

Step 1 Be sure SSH2 is enabled on this switch.

Step 2 Select **Show Cores** from the Admin menu.

You see the dialog box in Figure 64-4.

Figure 64-4 Show Cores Dialog

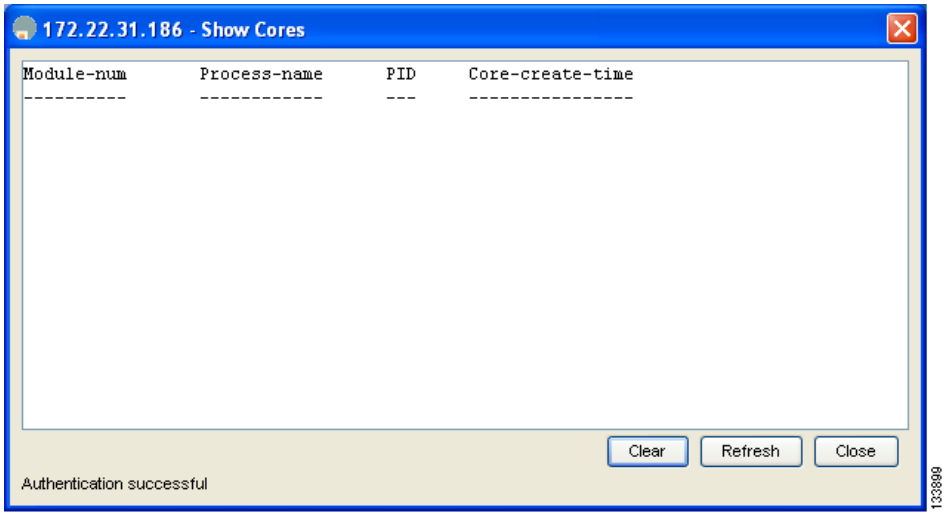

Step 3 Click **Clear** to clear the cores.

The software keeps the last few cores per service and per slot and clears all other cores present on the active supervisor module.

Online Health Management System

Note For information on most Online Health Management System procedures, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

The Online Health Management System (OHMS) is a hardware fault detection and recovery feature. It runs on all Cisco MDS switching, services, and supervisor modules and ensures the general health of any switch in the Cisco MDS 9000 Family. The OHMS monitors system hardware in the following ways:

- The OHMS component running on the active supervisor maintains control over all other OHMS components running on the other modules in the switch.
- The system health application running in the standby supervisor module only monitors the standby supervisor module—if that module is available in the HA standby mode. See the "HA Switchover Characteristics" section on page 15-2.

The OHMS runs multiple tests on each module at pre-configured intervals. The tests cover all major fault points, and isolate any failing component in the MDS switch. On detecting a fault, the OHMS does the following:

- Performs additional testing to isolate the faulty component
- Attempts to reconfigure the component by retrieving its configuration information from persistent storage
- If unable to recover, sends Call Home notifications, system messages and exception logs; and shuts down and discontinues testing the failed module or component (such as an interface)

If an error is detected on the active supervisor module and a standby supervisor module exists, the OHMS switches to the standby supervisor module and runs active supervisor tests on that module after it restarts. If a standby supervisor module does not exist in the switch, the OHMS reloads the switch.

Tests that can be performed using the OHMS include:

- Performing Internal Loopback Tests, page 64-6
- Performing External Loopback Tests, page 64-7

Performing Internal Loopback Tests

You can run manual loopback tests to identify hardware errors in the data path in the switching or services modules, and the control path in the supervisor modules. Internal loopback tests send and receive FC2 frames to/from the same ports and provide the round trip time taken in microseconds. These tests are available for Fibre Channel, IPS, and iSCSI interfaces.

Choose **Interface > Diagnostics > Internal** to perform an internal loopback test from Device Manager.

Performing External Loopback Tests

You can run manual loopback tests to identify hardware errors in the data path in the switching or services modules, and the control path in the supervisor modules. External loopback tests send and receive FC2 frames to/from the same port or between two ports.

You need to connect a cable (or a plug) to loop the Rx port to the Tx port before running the test. If you are testing to/from the same port, you need a special loop cable. If you are testing to/from different ports, you can use a regular cable. This test is only available for Fibre Channel interfaces.

Choose **Interface > Diagnostics > External** to perform an external loopback test from Device Manager.

Default Settings

Table 64-1 lists the default system health and log settings.

Table 64-1 Default System Health and Log Settings

Parameters	Default
Kernel core generation	One module.
System health	Enabled.
Loopback frequency	5 seconds.
Failure action	Enabled.

APPENDIX A

Cisco Fabric Manager Unsupported Feature List

This appendix contains a list of features and functions not supported by Cisco Fabric Manager or Device Manager. This list is organized according to the chapter in which the feature would be described if it were supported. (See Table A-1.) For documentation about these features, refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

Table A-1 Features Not Supported by Cisco Fabric Manager or Device Manager (contd.)

Part	Chapter/Category	Procedure
2 Cisco MDS SAN-OS Installation and Switch Management	Obtaining and Installing Licenses	Backing Up License Files Updating Licenses Moving Licences Between Switches
	Initial Configuration	Starting a Switch (Initial Setup) Configuring Console Settings Configuring COM1 and Modem Settings Adjusting for Daylight Savings Time Configuring the Initialization String Basic Switch Configuration Terminal Settings File System Commands Displaying File Contents
	Software Images	Manual Upgrade on a Dual Supervisor Switch Corrupted Bootflash Recovery

Part	Chapter/Category	Procedure
	Working with Configuration Files	Formatting External CompactFlash
		Compressing and Uncompressing Files
		Displaying the Last Lines in a File
		Executing Commands Specified in a Script
		Setting the Delay Time
		Displaying Configuration Files
		Unlocking the Startup Configuration File
		Accessing Remote File Systems
	Configuring High Availability	Copying Images to the Standby Supervisor
	Managing System Hardware	Clock Modules
	Managing Modules	Connecting to a Module
		Preserving Module Configuration
		Purging Module Configuration
		Reloading the Switch
		EPLD Configuration
		ASM-FSN Boot Image
		Configuring SSI Boot Image
		Managing ASM and SSM Modules
3 Switch Configuration	Configuring Interfaces	Displaying the ALPA Cache Contents
		Clearing the ALPA Cache
		N-Port Identifier Virtualization (NPIV)
	Scheduling Tasks	Schedule Configuration
4 Fabric Configuration	Inter-VSAN Routing Configuration	Inter-VSAN Routing (IVR) FICON Support
		IVR Service Groups
	Distributing Device Alias Services	Configuring DDAS
	Managing FLOGI, Name Server, FDMI, and RSCN Databases	Suppressing Domain Format SW-RSCNs
6 IP Services	Configuring FCIP	Displaying and Clearing ARP Caches
	Configuring the SAN Extension Tuner	Tuning Configuration
	Configuring IP Storage	IPS Module Core Dumps
8 Network and Switch Monitoring	Monitoring Network Traffic Using SPAN	Remote SPAN

Part	Chapter/Category	Procedure
10 Troubleshooting	Troubleshooting Your Fabric	Loop Monitoring Configuring CIM CFS for FC Timers Local Text Based Capture Capturing FC Analyzer Frames Locally Sending Captured FC Analyzer Frames to a Remote IP Address Clearing Configured FC Analyzer Information Displaying a List of Hosts Configured for Remote Capture Using Fabric Analyzer Display Filters
	Monitoring System Processes and Logs	Saving the Last Core to Flash Kernel Core Dumps System Health Initiation Loopback Test Configuration Frequency Hardware Failure Action Tests for a Specified Module Clearing Previous Error Reports Online Health Management System • Enabling and Disabling the OHMS • Enabling and Disabling Hardware Failure Action • Configuring Onboard Failure Logging • Clearing Previous Error Reports • Performing Tests for a Specified Module • Configuring Automatic Loopback Tests • Performing SERDES Loopback Tests

APPENDIX B

Interface Nonoperational Reason Codes

If the administrative state for an interface is up and the operational state is down, the reason code differs based on the nonoperational reason code as described in Table B-1.

Table B-1 Reason Codes for Nonoperational States

Reason Code	Description	Applicable Modes
Link failure or not connected	Physical layer link is not operational.	All
SFP not present	The small form-factor pluggable (SFP) hardware is not plugged in.	
Initializing	The physical layer link is operational and the protocol initialization is in progress.	
Reconfigure fabric in progress	The fabric is currently being reconfigured.	
Offline	Cisco MDS SAN-OS waits for the specified R_A_TOV time before retrying initialization.	
Inactive	The interface VSAN is deleted or is in a suspended state. To make the interface operational, assign that port to a configured and active VSAN.	
Hardware failure	A hardware failure is detected.	
Error disabled	Error conditions require administrative attention. Interfaces may be error-disabled for various reasons. For example: • Configuration failure. • Incompatible buffer-to-buffer credit configuration. To make the interface operational, you must first fix the error conditions causing this state; and next, administratively shut down or enable the interface.	

Table B-1 Reason Codes for Nonoperational States (continued)

Reason Code	Description	Applicable Modes
Isolation due to ELP failure	Port negotiation failed.	Only E ports and TE ports
Isolation due to ESC failure	Port negotiation failed.	
Isolation due to domain overlap	The Fibre Channel domains (fcdomain) overlap.	
Isolation due to domain ID assignment failure	The assigned domain ID is not valid.	
Isolation due to other side E port isolated	The E port at the other end of the link is isolated.	
Isolation due to invalid fabric reconfiguration	The port is isolated due to fabric reconfiguration.	
Isolation due to domain manager disabled	The fcdomain feature is disabled.	
Isolation due to zone merge failure	The zone merge operation failed.	
Isolation due to VSAN mismatch	The VSANs at both ends of an ISL are different.	
Nonparticipating	FL ports cannot participate in loop operations. It may happen if more than one FL port exists in the same loop, in which case all but one FL port in that loop automatically enters nonparticipating mode.	Only FL ports and TL ports
PortChannel administratively down	The interfaces belonging to the PortChannel are down.	Only PortChannel interfaces
Suspended due to incompatible speed	The interfaces belonging to the PortChannel have incompatible speeds.	
Suspended due to incompatible mode	The interfaces belonging to the PortChannel have incompatible modes.	
Suspended due to incompatible remote switch WWN	An improper connection is detected. All interfaces in a PortChannel must be connected to the same pair of switches.	

APPENDIX C

Managing Cisco FabricWare

The Cisco FabricWare software running on the MDS 9020 Switch offers Fibre Channel switching services that realize maximum performance. Cisco FabricWare provides networking features such as zoning, advanced security, nondisruptive software upgrades, diagnostics, a CLI with syntax resembling Cisco IOS, and standard interfaces for management applications.

This appendix contains the following sections:

Fibre Channel Support

Cisco FabricWare supports autoconfigured Fibre Channel ports capable of up to 4-Gbps bandwidth. Cisco FabricWare supports the following port types:

- E
- F
- FL
- Fx
- Auto

See the "About Interface Modes" section on page 18-3.

Cisco FabricWare supports Fabric Shortest Path First (FSPF) as the standard path selection protocol used by Fibre Channel fabrics. The FSPF feature is enabled by default on all Fibre Channel switches. Except in configurations that require special consideration, you do not need to configure any FSPF services. FSPF automatically calculates the best path between any two switches in a fabric.

Zone Configuration

Zoning enables you to set up access control between storage devices or user groups. If you have administrator privileges in your fabric, you can create zones to increase network security and to prevent data loss or corruption. Zoning is enforced by examining the source-destination ID field. Cisco FabricWare does not support QoS, broadcast, LUN, or read-only zones.

You can use the Fabric Manager zone configuration tool to manage zone sets, zones, and zone membership for switches running Cisco FabricWare. Cisco FabricWare supports zone membership by pWWN. See the "Configuring a Zone Using the Zone Configuration Tool" section on page 26-7.

Security

Cisco FabricWare supports the following security features:

- RADIUS
- SSH
- User-based roles
- IP access control lists

Cisco FabricWare can use the RADIUS protocol to communicate with remote AAA servers. RADIUS is a distributed client/server protocol that secures networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco MDS 9000 Family switches and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.

You can access the CLI using the console (serial connection), Telnet, or Secure Shell (SSH). For each management path (console or Telnet and SSH), you can configure one or more of the following security control options: **local**, **remote (RADIUS),** or **none**.

Using these access methods, you can configure the roles that each authenticated user receives when they access the switch. Cisco FabricWare supports two fixed roles: network administrator and network operator.

IP access lists (IP-ACLs) control management traffic over IP by regulating the traffic types that are allowed or denied to the switch. IP-ACLs can only be configured for the mgmt0 port.

Fabric Manager Server uses SNMPv1 and SNMPv2 to communicate with Cisco FabricWare.

Events

You can monitor fabric and switch status for Cisco FabricWare switches through either a syslog server or an SNMP trap receiver.

The syslog, or system message logging software, saves messages in a log file or directs the messages to other devices. This feature provides you with the following capabilities:

- Provides logging information for monitoring and troubleshooting
- Allows you to select the types of captured logging information
- Allows you to select the destination server to forward the captured logging information

By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. You can specify which system messages should be saved based on the type of facility and the severity level. You can access logged system messages using the CLI or by saving them to a properly configured system message logging server.

You can configure the Cisco MDS 9020 Switch using the CLI to send notifications to SNMP managers when particular events occur. You can send these notifications as traps.

Managing Cisco FabricWare with Fabric Manager

Fabric Manager Release 2.1(2) or later supports switches running Cisco FabricWare. Table C-1 shows the supported features and where to find more information on that feature.

Note You cannot configure port speeds using Fabric Manager or Device Manager on the Cisco MDS 9020 Fabric Switch.

Note You cannot configure users using Fabric Manager or Device Manger on the Cisco MDS 9020 Fabric Switch.

Table C-1 FabricWare Features in Fabric Manager

Feature	FabricWare Capabilities	Section
Zones	Zone configuration Zone membership by pWWN No Cisco FabricWare support for QoS, broadcast, LUN, or read-only zones	"Configuring a Zone Using the Zone Configuration Tool" section on page 26-7 "Adding Zone Members" section on page 26-9 "About Zoning" section on page 26-2
Interfaces	1/2/4 Fibre Channel autonegotiating ports	"Fibre Channel Interfaces" section on page 18-1
SNMP	SNMPv1 and SNMPv2c	"SNMP Version 1 and Version 2c" section on page 34-2
Software images	Automated upgrades Manual upgrades	"Using the Software Install Wizard" section on page 13-7 "Software Upgrade Methods" section on page 13-4
FLOGI, name server, FDMI, and RSCN	Displaying FLOGI details Registering name server proxies Displaying FDMI RSCN statistics	Refer to the *Cisco MDS 9020 Switch Configuration Guide and Command Reference*.

Feature	FabricWare Capabilities	Section
Security	Configuring RADIUS Configuring server groups Configuring role-based authorization Configuring user accounts Configuring SSH services	"Configuring a RADIUS Server" section on page 35-10 "Configuring Server Groups" section on page 35-20 "Role-Based Authorization" section on page 33-1 "Configuring Users" section on page 33-12 "Enabling SSH or Telnet Service" section on page 33-17
Fibre Channel routing	FSPF global configuration FSPF interface configuration	Refer to the *Cisco MDS 9020 Switch Configuration Guide and Command Reference*.
IP services	IP access control lists on mgmt0	"Creating IPv4-ACLs in Device Manager" section on page 36-6
System messages	System message logging configuration	"Viewing Logs from Device Manager" section on page 53-5
Advanced configuration	FC timer	"Fibre Channel Time Out Values" section on page 32-9

APPENDIX D

Configuration Limits for Cisco MDS SAN-OS Release 3.x

The features supported by Cisco MDS SAN-OS have maximum configuration limits. For some of the features, we have verified configurations that support limits less that the maximum. Table D-1 lists the Cisco verified limits and maximum limits for switches running Cisco MDS SAN-OS Release 3.x.

Table D-1 Cisco MDS SAN-OS Release 3.x Configuration Limits

Feature	Verified Limit	Maximum Limit
VSANs	80 VSANs per physical fabric.	4000 VSANs per physical fabric.
Switches in a single MDS physical fabric or VSAN	40 switches.	239 switches.
Switches in a mixed or open physical fabric or VSAN	32 switches.	239 switches.
Domains per VSAN	40 domains.	239 domains.
Zone members	16,000 zone members per physical fabric (includes all VSANs).	20,000 zone members per Physical Fabric (includes all VSANs).
Zones	8000 zones per switch (includes all VSANs).	8000 zones per switch (includes all VSANs).
Zone sets	500 zone sets per switch (includes all VSANs).	1000 zone sets per switch (includes all VSANs).
Supported hops for all major storage, server, and HBA vendors	7 hops (diameter of the SAN fabric).	12 hops.
IVR zone members	4000 IVR zone members per physical fabric.	10,000 IVR zone members per physical fabric.
IVR zones	1500 IVR zones per physical fabric.	2000 IVR zones per physical fabric.
IVR zone sets	32 IVR zone sets per physical fabric.	32 IVR zone sets per physical fabric.
IVR service groups	16 service groups per physical fabric.	16 service groups per physical fabric.

Table D-1 Cisco MDS SAN-OS Release 3.x Configuration Limits (continued)

Feature	Verified Limit	Maximum Limit
ISL instances per switch [1]	Up to 200 ISLs, each with 16 VSANs, for a total of 3200 port-VSAN instances. You can configure more than 200 ISLs with fewer than 16 VSANs, or fewer than 200 ISLs with more than 16 VSANs, within the total ports per VSAN instance limit of 3200.	Up to 200 ISLs, each with 16 VSANs, for a total of 3200 port-VSAN instances. You can configure more than 200 ISLs with fewer than 16 VSANs, or fewer than 200 ISLs with more than 16 VSANs, within the total ports per VSAN instance limit of 3200.
IP ports per switch	No limits.	No limits.
Fibre Channel modules vs. IPS modules per switch	No limits.	No limits.
iSCSI and iSLB sessions per IP port	500 sessions.	500 sessions.
iSCSI and iSLB sessions per switch	5000 sessions.	5000 sessions.
iSCSI and iSLB initiators supported in physical fabric	2000 initiators.	2000 initiators.
iSCSI and iSLB targets per physical fabric (virtual and initiator targets)	6000 targets.	6000 targets.

1. This is the number of trunking-enabled ISL ports multiplied by the number of VSANs in the switch.

INDEX

Symbols

Numerics

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

Z