Cisco MDS 9000 Fabric Manager Switch Configuration Guide

March 2004

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-7753-01

Cisco MDS 9000 Fabric Manager Switch Configuration Guide

CONTENTS

Preface xxvii

Audience xxvii

Organization xxvii

Conventions xxx

Obtaining Documentation xxxi

Cisco.com xxxi

Documentation DVD xxxi

Ordering Documentation xxxi

Documentation Feedback xxxii

Cisco Product Security Overview xxxii

Reporting Security Problems in Cisco Products xxxii

Obtaining Technical Assistance xxxiii

Cisco Technical Support Website xxxiii

Submitting a Service Request xxxiii

Definitions of Service Request Severity xxxiv

Obtaining Additional Publications and Information xxxiv

New and Changed Information xxxvii

CHAPTER 1

Product Overview 1-1

Hardware Overview 1-1

Cisco MDS 9216 Fabric Switch 1-2

Cisco MDS 9500 Modular Directors 1-2

Cisco MDS 9100 Series Fixed Configuration Fabric Switches 1-3

Software Features 1-4

Licensing 1-4

High Availability 1-4

Switch Reliability 1-4

Virtual SANs 1-5

Intelligent Zoning 1-5

Inter-VSAN Routing 1-5

Trunking 1-6

PortChannels 1-6

IP Services 1-6

IP Storage 1-7
Call Home 1-7
QoS and Congestion Control 1-7
SPAN and RSPAN 1-8
Switch Management Features 1-8
Redundant Supervisor Module Management 1-8
Fabric Management 1-9
Security Management 1-9
Tools for Software Configuration 1-10
CLI 1-10
Cisco MDS 9000 Fabric Manager 1-11

CHAPTER 2 **Getting Started with Cisco Fabric Manager** 2-1
Managing Cisco MDS 9000 Switches 2-2
Storage Management Solutions Architecture 2-3
In-Band Management and Out-of-Band Management 2-4
MGMT0 2-4
IPFC 2-4
Installing the Applications 2-5
Launching the Applications 2-6
Using the Management Services Wizard 2-7
A Note on Ports 2-7

CHAPTER 3 **Overview of Fabric Manager** 3-1
Launching Cisco Fabric Manager 3-1
Using Fabric Manager 3-2
Menu Bar, Toolbars, and Status Bar 3-3
Logical/Physical Pane 3-3
Information Pane 3-4
Map Pane 3-4
Discovering and Viewing the Network Fabric 3-7
Controlling Administrator Access with Users and Roles 3-7
Modifying Device Grouping 3-7
Setting Fabric Manager Preferences 3-8
Viewing Reports in Fabric Manager 3-9
Using Device Manager 3-10
Launching Device Manager from Fabric Manager 3-10
Using Summary View 3-11

Comparing Device Manager to Fabric Manager 3-12
Performing Device Management 3-13
Managing Ports 3-13
Setting Device Manager Preferences 3-14
Using Performance Manager 3-14
Performance Manager Architecture 3-14
Creating a PM Configuration File 3-14
Collecting the Data 3-15
Presenting the Collected Data 3-15
Exporting and Importing Data 3-16
Integration with Cisco Traffic Analyzer 3-16
Configuring PM for Use with Cisco Traffic Analyzer 3-16
Stopping Data Collection 3-19
Exporting Data Collection to XML Files 3-19
Removing Data Collection Files from the List 3-19

CHAPTER 4 **Before You Begin** 4-1
About Flash Devices 4-1
Internal bootflash: 4-2
External CompactFlash (Slot0) 4-2
Switch Roles 4-2
Using Valid Formats and Ranges 4-2

CHAPTER 5 **Obtaining and Installing Licenses** 5-1
License Terminology 5-1
Licensing Model 5-2
Licensing High Availability 5-4
Options to Install a License 5-4
Obtaining a Factory-Installed License 5-4
Performing a Manual Installation 5-5
Obtaining License Key Files 5-5
Installing Licenses 5-6
Installing Licenses Using Fabric Manager License Wizard 5-6
Installing Licenses Using Device Manager 5-8
Viewing License Information in Fabric Manager 5-8
Viewing License Information in Device Manager 5-9
Removing Licenses 5-9

Updating Licenses 5-10
License Expiry Alerts 5-10
Moving Licenses Between Switches 5-11

CHAPTER 6 **Initial Configuration** 6-1
NTP Configuration 6-1
NTP Configuration Guidelines 6-2
Display General NTP Statistics for a Switch 6-3
Create an NTP Server or Peer 6-3
Edit an NTP Server or Peer Configuration 6-4
Delete an NTP Server or Peer 6-4

CHAPTER 7 **Configuring High Availability** 7-1
About High Availability 7-1
Switchover Mechanisms 7-2
HA Switchover 7-2
Process Restartability 7-2
Synchronizing Supervisor Modules 7-2
HA Redundancy States 7-2

CHAPTER 8 **Software Images** 8-1
About Software Images 8-1
Essential Upgrade Prerequisites 8-2
Using the Software Install Wizard 8-3
Maintaining Supervisor Modules 8-4
Standby Supervisor Boot Variable Version 8-4
Standby Supervisor Boot Alert 8-5
Replacing Modules 8-5
Recovering a Corrupted Bootflash 8-5
Default Factory Settings 8-5

CHAPTER 9 **Managing Modules** 9-1
About Modules 9-1
Supervisor Modules 9-1
Switching Modules 9-2
Viewing the State of a Module 9-2
Identifying Module LEDs 9-3

Configuring EPLDs 9-5
Default Supervisor Module Settings 9-6

CHAPTER 10 **Managing System Hardware** 10-1
Configuring Power Supplies 10-1
Guidelines for Power Supplies with Different Capacities 10-1
Guidelines for Power Supplies with Different Capacities 10-3
Managing Power Supplies 10-4
Displaying Module Temperature 10-5
Monitoring Fan Modules 10-6
Monitoring Clock Modules 10-6
Viewing System Attributes 10-6
Viewing Running Processes 10-6
Viewing Flash File Information 10-7
Managing Inventory Information 10-7
Managing Module Attributes 10-7
10-7

CHAPTER 11 **Configuring and Managing VSANs** 11-1
How VSANs Work 11-2
VSANs Versus Zones 11-4
Default and Isolated VSANs 11-5
Default VSANs 11-5
Isolated VSANs 11-5
VSAN Membership 11-6
VSAN Attributes 11-6
Operational State of a VSAN 11-6
Adding and Configuring VSANs 11-7
Deleting VSANs 11-7
Default Settings 11-8

CHAPTER 12 **Configuring Interfaces** 12-1
Configuring Fibre Channel Interfaces 12-1
About Interface Modes 12-2
E Port 12-2
F Port 12-2
FL Port 12-3

TL Port 12-3
TE Port 12-3
SD Port 12-3
ST Port 12-4
Fx Port 12-4
B Port 12-4
Auto Mode 12-4
About Interface States 12-4
Administrative States 12-4
Operational States 12-5
Reason Codes 12-5
Configuring TL Port ALPA Caches 12-8
Configuring Buffer-to-Buffer Credits 12-8
Configuring Performance Buffers 12-8
Configuring the Beacon Mode 12-9
Identifying the Beacon LEDs 12-9
Configuring Switch Port Defaults 12-10
Default Settings 12-10
Configuring VSAN Interfaces 12-10
Configuring Gigabit Ethernet Interfaces 12-11
Enabling or Disabling Ports 12-11
Managing Interface Attributes for Ports 12-11

CHAPTER 13 **Configuring Trunking** 13-1
About Trunking 13-1
About Trunking Protocol 13-2
Configuring Trunk Modes 13-2
Configuring Trunk-Allowed VSAN List 13-2
Trunking Configuration Guidelines 13-4
Default Settings 13-5

CHAPTER 14 **Configuring PortChannels** 14-1
PortChannel Examples 14-1
Configuring 32-port Switching Modules and Host-Optimized Ports 14-2
Managing Physical Attributes for a Port 14-2
Viewing Port Capability Attributes 14-3
About PortChanneling and Trunking 14-3
Managing PortChannel General Attributes 14-4

Managing PortChannel Interface Attributes 14-4
Quiescing/Disabling Port Channel Members 14-4
About Load Balancing 14-5
Considerations for PortChannel Configurations 14-6
Error Detection 14-6
Default Settings 14-7

CHAPTER 15 **Configuring and Managing Zones** 15-1
Zoning Features 15-2
Zoning Example 15-3
Configuring a Zone 15-4
Creating Zones 15-4
Creating Additional Zones 15-4
Cloning Zones 15-5
Adding Zone Members 15-5
Displaying Port Membership Information 15-6
Viewing Zone Statistics 15-6
Deleting Zones and Members 15-6
Configuring Aliases 15-6
Creating Zones with Aliases 15-7
Viewing Aliases 15-7
Zone Sets 15-7
Active and Full Zone Set Considerations 15-8
Distributing Zone Sets 15-10
Copying Zone Sets 15-10
Creating Zone Sets 15-10
Creating Additional Zone Sets 15-10
Cloning Zone Sets 15-11
Adding Zones to a Zone Set 15-11
Activating or Enforcing Zone Sets 15-11
Deactivating Zone Sets 15-12
Importing Active Zone Sets 15-12
Exporting Active Zone Sets 15-12
Deleting Zone Sets or Members 15-12
Clearing the Zone Database 15-13
Recovering a Full Zone Database 15-13
Performing Zone Merge Analysis 15-13
Zone Enforcement 15-14

The Default Zone 15-14
Setting Default Zone Policy 15-15
Changing the Default Zone Policy 15-15
Recovering from Link Isolation 15-15
LUN Zoning 15-16
Assigning LUNs to Storage Subsystems 15-17
Read-Only Zoning 15-17
Guidelines to Configure Read-Only Zones 15-17
Default Settings 15-18
Migrating a Non-MDS Database 15-18
Using the Zone Wizard 15-18

CHAPTER 16 **Configuring Inter-VSAN Routing** 16-1
About IVR 16-1
IVR Features 16-2
IVR Terminology 16-2
IVR Guidelines 16-3
Domain ID Guidelines 16-3
Transit VSANs Guidelines 16-3
Border Switch Guidelines 16-3
Configuring IVR 16-4
Unique Domain ID Configuration Options 16-4
Enabling IVR 16-4
Configuring an IVR Topology 16-4
Creating an IVR Topology 16-4
Creating IVR Zones and Zone Sets 16-5
Creating Additional IVR Zones and Zone Sets 16-5
Activating IVR Zone Sets 16-6
Deactivating IVR Zone Sets 16-6
Recovering an IVR Full Zone Database 16-6
Recovering an IVR Full Topology 16-7
IVR Interoperability 16-7
IVR Using LUN Zoning or Read-Only Zoning 16-7
Creating IVZs and IVZSs 16-7
Zones versus IVZs 16-8
Automatic IVZ Creation 16-8
Configuring and Activating IVZs and IVZSs 16-9
Using the force Option 16-9

Clearing the IVZ Database 16-9
Using the Zone Wizard 16-10

CHAPTER 17 **Managing FLOGI, Name Server, FDMI, and RSCN Databases** 17-1
Displaying FLOGI Details 17-1
Configuring the Name Server Proxy Feature 17-2
Displaying Name Server Database Entries 17-2
Displaying FDMI 17-2
Displaying RSCN Information 17-2
Sending RSCNs 17-3
Viewing General Attributes for the Name Server 17-3
Viewing Advanced Attributes for the Name Server 17-3
Proxy Ports for the Name Server 17-3
Viewing Name Server Statistics 17-4
Viewing RSCN Nx Registrations 17-4
Viewing RSCN Statistics 17-4
Viewing FLOGI Attributes 17-4
Viewing Port ELP Attributes 17-5
Viewing Trunk Configuration 17-5

CHAPTER 18 **Configuring Switch Security** 18-1
Switch Management Security 18-2
SNMP Security 18-2
CLI Security 18-2
Switch AAA Functionalities 18-2
Authentication 18-3
Authorization 18-3
Accounting 18-3
Remote Authentication by AAA Servers 18-3
Remote Authentication Guidelines 18-3
Server Groups 18-4
AAA Service Configuration Options 18-4
Configuring RADIUS 18-4
About RADIUS 18-5
Configuring RADIUS Authentication 18-5
Configuring RADIUS Servers 18-5
Setting the RADIUS Server Address 18-5

Setting the RADIUS Preshared Key 18-6
Setting Iterations of the RADIUS Server 18-6
Defining Vendor-Specific Attributes 18-6
VSA Format 18-6
Configuring TACACS+ 18-7
About TACACS+ 18-7
Advantages of TACACS+ 18-7
Enabling TACACS+ 18-8
Setting the TACACS+ Server Address 18-8
Setting the Secret Key 18-8
Setting the Timeout Value 18-8
Defining Custom Attributes for Roles 18-8
Configuring Server Groups 18-9
Local AAA 18-9
Authentication and Authorization Process 18-9
Configuring Role-Based CLI Authorization 18-11
Configuring Rules and Features for Each Role 18-11
Configuring the VSAN Policy 18-12
Recovering Administrator Password 18-12
Configuring SSH Services 18-12
Enabling SSH Service 18-13
Generating an SSH Host Key Pair 18-13
Using the force Option 18-13
About SNMP Security 18-13
SNMP Version 1 and Version 2c 18-14
Adding a Community String 18-14
Deleting a Community String 18-14
SNMP Version 3 18-14
Adding SNMP Users 18-15
Deleting SNMP Users 18-15
Configuring and Creating SNMP User Roles 18-15
Viewing SNMP Community and User Information 18-16
Group-Based SNMP Access 18-16
Configuring Common Roles 18-16
Creating and Modifying Users 18-17
Creating Common Roles 18-18
Editing Common Role Rules (Device Manager Only) 18-19
Deleting Common Roles 18-19
Assigning Users to Roles 18-19

Default Security Settings 18-20
Restricting Switch Access 18-21

CHAPTER 19 **Configuring Fabric Security** 19-1
About Fabric Authentication 19-1
About DHCHAP 19-2
DHCHAP Compatibility with Existing MDS Features 19-2
Configuring DHCHAP Authentication 19-3
Enabling DHCHAP 19-3
Configuring DHCHAP Authentication Modes 19-3
Configuring the DHCHAP Hash Algorithm 19-4
Configuring DHCHAP Groups 19-4
Configuring DHCHAP Passwords 19-4
Configuring Passwords for Other Devices 19-5
Configuring the DHCHAP Timeout Value 19-5
Default Fabric Security Settings 19-5

CHAPTER 20 **Configuring Port Security** 20-1
Port Security Features 20-1
Enforcing Port Security 20-1
Configuring a Port Binding 20-2
Copying an Active Configuration to the Running Configuration 20-2
Deleting a Port Binding 20-3
About Auto-Learn 20-3
Activating Port Security 20-3
Activating a Port Binding 20-3
Displaying Activated Port Bindings 20-4
Configuring Auto-Learning 20-4
Authorization Scenario 20-5
Turning Auto-Learning On or Off 20-6
Manually Configuring Port Security 20-7
Identifying WWNs to Configure Port Security 20-7
Securing Authorized Ports 20-7
Activating the Port Security Database 20-7
Forcing Port Security Activation 20-8
Reactivating the Database 20-8
Database Scenarios 20-8

Displaying Port Security Statistics 20-9
Displaying Port Security Violations 20-9
Default Port Security Settings 20-9

CHAPTER 21 **Configuring Fibre Channel Routing Services and Protocols** 21-1
FSPF Features 21-2
FSPF Examples 21-2
Fault Tolerant Fabric 21-2
Redundant Links 21-3
Fail-over Scenarios for PortChannels and FSPF Links 21-3
Configuring FSPF Globally 21-4
Managing FSPF General Attributes 21-4
Disabling FSPF Routing Protocols 21-4
Link State Record Defaults 21-4
Viewing Link State Records 21-5
Viewing FSPF Links 21-5
Configuring FSPF for a Specific Interface 21-5
Configuring FSPF Interfaces 21-6
Computing Route Cost 21-6
Specifying Hello Time Intervals 21-6
Specifying Dead Intervals 21-6
Disabling FSPF for Specific Interfaces 21-6
Retransmitting Intervals 21-6
Viewing FSPF Interface Statistics 21-7
Configuring Fibre Channel Routes 21-7
Configuring Fibre Channel Route Flows 21-7
Broadcast Routing 21-8
In-Order Delivery 21-8
Reordering Network Frames 21-8
Reordering PortChannel Frames 21-9
Enabling In-Order Delivery 21-10
Configuring Flow Statistics 21-10
Viewing FSPF Statistics 21-10
Default Settings 21-10

CHAPTER 22 **Configuring IP Services** 22-1
Traffic Management Services 22-2
Configuring the Ethernet Management Port 22-2

Configuring the Default Gateway 22-3
Configuring the Default Network 22-4
Configuring an IP Route 22-4
IP Access Control Lists 22-4
IP-ACL Configuration Guidelines 22-4
Creating IP-ACLs 22-5
Adding Entries to an Existing IP-ACL 22-5
Comparing Ports 22-5
Applying IP-ACLs 22-7
Configuring IPFC 22-8
Configuring Overlay VSANs 22-8
Configuring Multiple VSANs 22-9
Managing IPFC Connectivity with Multiple VSANs 22-10
Configuring VRRP 22-10
VRRP Features 22-11
VRRP Functionality 22-11
Creating or Removing a Virtual Router 22-12
Enabling a Virtual Router 22-12
Adding an IP Address for a Virtual Router 22-12
Viewing IP Address Information 22-12
Managing IP Addresses for VRRP 22-13
Setting Priority for the Virtual Router 22-13
Setting the Time Interval for the Advertisement Packet 22-13
Preempting the Master Virtual Router 22-13
Configuring Authentication for the Virtual Router 22-13
Setting the Priority Based on Interface State 22-14
Configuring VRRP Operations Attributes 22-14
Default Settings 22-14
Enabling or Disabling IP Forwarding 22-15
Viewing Information and Statistics 22-15
Viewing VRRP Statistics 22-15
Viewing TCP Information and Statistics 22-15
Viewing UDP Information and Statistics 22-15
Viewing IP Statistics 22-16
Viewing ICMP Statistics 22-16

CHAPTER 23 **Configuring FICON** 23-1
About FICON 23-2

MDS-Specific FICON Advantages 23-3
Fabric-Optimization with VSANs 23-3
FCIP Support 23-4
PortChannel Support 23-5
VSANs for FICON and FCP Intermixing 23-5
MDS-Supported FICON Features 23-5
FICON Port Numbering 23-6
Port Addresses 23-8
Implemented and Unimplemented Port Addresses 23-9
Installed and Uninstalled Ports 23-9
FCIP Port Number 23-9
Port Numbering Summary 23-10
FCIP and PortChannel Port Numbers 23-10
FC ID Allocation 23-11
FICON Cascading 23-11
MDS FICON Prerequisites 23-11
Enabling FICON 23-12
Effects of Enabling FICON 23-12
Creating FICON VSANs (enabling FICON) Using Fabric Manager 23-12
Creating FICON VSANs (enabling FICON) Using Device Manager 23-13
Deleting FICON VSANs (Disabling FICON) 23-13
Viewing FICON Director History 23-14
Configuring Code Page 23-14
Configuring the FC ID Last Byte 23-14
Automatically Saving the Running Configuration 23-14
Binding Port Numbers to PortChannels 23-15
Binding Port Numbers to FCIP Interfaces 23-15
Configuring FICON Ports 23-15
FICON Information Refresh Note 23-15
Blocking Ports 23-15
Prohibiting Ports 23-16
Entering FICON Port Configuration Information 23-16
Viewing FICON Port Attributes 23-17
FICON Configuration Files 23-17
Accessing FICON Configuration Files 23-18
Editing FICON Configuration Files 23-18
Creating FICON Files 23-18
Deleting FICON Files 23-19

Copying FICON Files 23-19
Port Swapping 23-19
Port Swapping Guidelines 23-20
Swapping FICON Ports 23-20
Clearing FICON Device Allegiance 23-21
CUP Inband Management 23-21
Configuring Fabric Binding 23-21
Port Security versus Fabric Binding 23-22
Enforcing Fabric Binding 23-22
Enabling Fabric Binding 23-23
Configuring a List of sWWNs 23-23
Activating Fabric Binding 23-23
Forcing Fabric Binding Activation 23-23
Activating Fabric Binding 23-24
Deactivating Fabric Binding 23-24
Fabric Binding CopyActive to Config 23-24
Creating a Fabric Binding Configuration 23-25
Deleting a Fabric Binding Configuration 23-25
Viewing Fabric Binding Active Database 23-25
Viewing Fabric Binding Violations 23-25
Clearing Fabric Binding Statistics 23-26
Viewing EFMD Statistics 23-26
Displaying RLIR Information 23-26

CHAPTER 24 **Configuring IP Storage** 24-1
IP Storage Services Module 24-1
Configuring Gigabit Ethernet Interfaces 24-2
About Gigabit Ethernet Interfaces 24-3
Basic Gigabit Ethernet Configuration 24-3
About VLANs for Gigabit Ethernet 24-3
Verifying Gigabit Ethernet Connectivity 24-4
Gigabit Ethernet High Availability 24-4
Configuring VRRP 24-4
Configuring Ethernet PortChannels 24-5
Configuring CDP 24-6
IPS Core Dumps 24-6
Configuring FCIP 24-7

About FCIP 24-7
FCIP and VE Ports 24-7
FCIP Link 24-8
FCIP Profiles 24-8
FCIP Interface 24-9
Enabling FCIP 24-9
Basic FCIP Configuration 24-9
Creating FCIP Profiles 24-10
Creating FCIP Links 24-10
Creating FCIP Tunnels with Device Manager 24-11
Assigning FCIP Profiles 24-11
Creating Tunnels 24-11
Verifying Interfaces 24-12
Verifying Extended Link Protocols (ELP) 24-12
Checking Trunk Status 24-12
Checking for Interface Errors 24-13
Creating FCIP Tunnels with Fabric Manager 24-13
Advanced FCIP Profile Configuration 24-13
Configuring TCP Listener Ports 24-14
Configuring TCP Parameters 24-14
Advanced FCIP Interface Configuration 24-16
Configuring Peers 24-16
Configuring Active Connection 24-17
Configuring the Number of TCP Connections 24-17
Enabling Time Stamps 24-17
B Port Interoperability Mode 24-17
E Port Configurations 24-19
Configuring FCIP Write Acceleration 24-20
Enabling FCIP Compression 24-21
Fibre Channel PortChannels 24-22
FSPF 24-22
VRRP 24-23
Ethernet PortChannels 24-23
Ethernet PortChannels and Fibre Channel PortChannels 24-24
Configuring iSCSI 24-24
About iSCSI 24-25
Enabling iSCSI 24-26
Using the iSCSI Wizard 24-26
Routing iSCSI Requests and Responses 24-27
Presenting Fibre Channel Targets as iSCSI Targets 24-27

Dynamic Importing 24-28
Static Importing 24-28
iSCSI Virtual Target Configuration Examples 24-31
Presenting iSCSI Hosts as Virtual Fibre Channel Hosts 24-32
Dynamic Mapping 24-33
Static Mapping 24-33
Making the Dynamic Initiator WWN Mapping Static 24-34
Assigning VSAN Membership to iSCSI Hosts 24-34
Assigning VSANs to a iSCSI Interface 24-34
Configuring iSCSI Proxy Initiators 24-34
Configuring the iSCSI Proxy Initiator 24-35
Access Control in iSCSI 24-36
Fibre Channel Zoning-Based Access Control 24-36
iSCSI-Based Access Control 24-36
Enforcing Access Control 24-37
iSCSI User Authentication 24-37
Configuring an Authentication Mechanism 24-37
Configuring an iSCSI RADIUS Server 24-38
Advanced iSCSI Configuration 24-38
Setting the QOS Values 24-39
iSCSI Forwarding Mode 24-39
iSCSI High Availability 24-39
Multiple IPS Ports Connected to the Same IP Network 24-39
VRRP-Based High Availability 24-40
Ethernet PortChannel-Based High Availability 24-41
iSCSI Authentication Setup Guidelines 24-42
Configuring Storage Name Services 24-42
Creating iSNS Profiles and Tagging Profiles 24-42
Creating an iSNS Profile 24-43
Modifying an iSNS Profile 24-43
Default IP Storage Settings 24-44
Using the IP Filter Wizard 24-45
Creating IP Profiles 24-45
Adding IP Filters to Profiles 24-46
Associating IP Profiles to Interfaces 24-46
Deleting IP Profiles 24-47
Deleting IP Filters 24-47

CHAPTER 25 **Configuring Call Home** 25-1

Call Home Features 25-2
Call Home Configuration Process 25-2
Cisco AutoNotify 25-2
Assigning Contact Information 25-3
Configuring Destination Profiles 25-3
Configuring Alert Groups 25-3
Configuring Message Levels 25-4
Configuring E-Mail Options 25-4
Enabling or Disabling Call Home 25-4
Default Settings 25-4
Event Triggers 25-4
Call Home Message Severity Levels 25-7
Message Contents 25-8
Call Home Configuration Overview 25-11
Configuring Call Home Attributes 25-12
Configuring Call Home Destination Attributes 25-12
Configuring Call Home E-Mail Addresses 25-12
Configuring Call Home Alerts 25-13
Configuring Call Home Profiles 25-13

CHAPTER 26 **Configuring Domain Parameters** 26-1

About fcdomain Phases 26-2
Restarting the Domain 26-3
Performing a Domain Restart 26-3
Configuring the Domain 26-3
Configuring Domain Attributes 26-4
Managing Running Attributes for Domains 26-5
Viewing Domain Information 26-5
Viewing Domain Manager Statistics 26-5
Configuring Domain Interfaces 26-5
Viewing Domain Areas 26-6
Viewing Domain Area Ports 26-6
Specifying a Preferred or Static Domain ID 26-6
Setting Switch Priority 26-6
Configuring Allowed Domain ID Lists 26-6

Merging Stable Fabrics 26-7
Assigning Contiguous Domains 26-7
Disabling the fcdomain Feature 26-7
Setting the Fabric Name 26-8
Stopping Incoming RCFs 26-8
Configuring Persistent FC IDs 26-8
Creating a Persistent FC ID 26-9
Deleting a Persistent FC ID 26-9
Configuring Persistent FC IDs Manually 26-10
Configuring Unique Area FC IDs for Some HBAs 26-10
Enabling Persistent FC IDs 26-10
Purging Persistent FC IDs 26-11
Default Settings 26-12

CHAPTER 27 **Configuring Traffic Management** 27-1
FCC 27-1
FCC Process 27-2
Enabling FCC 27-2
QoS 27-2
Control Traffic 27-3
Disabling Control Traffic 27-3
Data Traffic 27-3
Configuring Data Traffic 27-4
Enabling QoS for Data Traffic 27-5
Creating Class Maps 27-5
Defining Service Policies 27-5
Applying a Service Policy 27-6
Scheduling Traffic 27-6
Ingress Port Rate Limiting 27-6
Default Settings 27-6

CHAPTER 28 **Configuring System Message Logging** 28-1
About System Message Logging 28-1
Configuring System Message Logging 28-3
Enabling Message Logging 28-4
Configuring Console Severity Level 28-4
Configuring Module Logging 28-4
Configuring Log Files 28-4

Configuring the Syslog Daemon 28-4
Outgoing Syslog Server Logging Facilities 28-5
Configuring Syslog Servers 28-6
Configuring Syslog Attributes 28-6
Configuring Syslog Priorities 28-6
Default Settings 28-6
About SNMP Events 28-7
Viewing the Events Log 28-7
Configuring Event Destinations 28-7
Configuring Event Security 28-8
Configuring Event Filters 28-8
About RMON Facilities 28-8
Enabling RMON Alarms by Port 28-8
Enabling RMON Alarms for VSANs 28-9
Enabling RMON Alarms for Physical Components 28-9
Configuring RMON Controls 28-10
Managing RMON Alarms 28-10
Managing RMON Event Severity Levels 28-10
Viewing the RMON Log 28-11

CHAPTER 29 **Discovering SCSI Targets** 29-1
About SCSI LUN Discovery 29-1
Initiating Customized Discovery 29-2
Authenticating iSCSI Targets 29-2
Specifying Targets 29-2
Specifying LUN Mappings 29-3
Viewing iSCSI Statistics 29-3
Viewing iSCSI Sessions 29-3
Viewing Session Statistics 29-4
Creating an iSCSI Initiator 29-4
Creating an iSCSI Virtual Target 29-4
Using the iSCSI Wizard 29-5

CHAPTER 30 **Monitoring Network Traffic Using SPAN** 30-1
About SPAN 30-1
SPAN Sources 30-2
IPS Source Ports 30-3
CSM Source Ports 30-3

Allowed Source Interface Types 30-3
VSAN as a SPAN Source 30-4
Guidelines to Configure VSANs as a Source 30-4
SPAN Sessions 30-5
Creating SPAN Sessions 30-5
Editing SPAN Sources 30-5
Deleting SPAN Sessions 30-6
Specifying Filters 30-6
Guidelines to Specifying Filters 30-6
SD Port Characteristics 30-7
Guidelines to Configure SPAN 30-7
Monitoring Traffic Using Fibre Channel Analyzers 30-7
Without SPAN 30-7
Using SPAN 30-8
Configuring Analyzers Using SPAN 30-9
Using a Single SD Port to Monitor Traffic 30-9
Default SPAN Settings 30-10
Remote SPAN 30-10
Advantages to Using RSPAN 30-11
FC and RSPAN Tunnels 30-11
Guidelines to Configure RSPAN 30-12
ST Port Characteristics 30-12
Configuring RSPAN 30-13
Configuration in the Source Switch 30-13
Configuration in All Intermediate Switches 30-14
Configuration in the Destination Switch 30-14
Configuring An Explicit Path 30-15
Monitoring RSPAN Traffic 30-16
Sample Scenarios 30-17
Single Source with One RSPAN Tunnel 30-17
Single Source with Multiple RSPAN Tunnels 30-17
Multiple Sources with Multiple RSPAN Tunnels 30-18

CHAPTER 31 **Advanced Features and Concepts** 31-1
Configuring FC Timers 31-1
Configuring Timers Per-VSAN 31-2
Configuring a Fabric Analyzer 31-2
About the Cisco Fabric Analyzer 31-3

Local Text-Based Capture 31-3
Remote Capture Daemon 31-4
GUI-Based Client 31-4
Configuring the Cisco Fabric Analyzer 31-4
Configuring World Wide Names 31-5
Allocating Flat FC IDs 31-5
Enabling Loop Monitoring 31-6
Configuring the Switch for Interoperability 31-6
Configuring Interoperability 31-8
Using the show tech-support Command 31-8
Managing World Wide Names 31-9
Configuring Timers 31-9

CHAPTER 32 **Configuring Fabric Configuration Servers** 32-1
About FCS 32-1
Significance of FCS 32-2

CHAPTER 33 **Monitoring System Processes and Logs** 33-1
Configuring Kernel Core Dumps 33-1

CHAPTER 34 **Troubleshooting the Fabric** 34-1
Analyzing Switch Device Health 34-1
Analyzing End-to-End Connectivity 34-2
Analyzing Switch Fabric Configuration 34-2
Analyzing the Results of Merging Zones 34-3
Issuing the Show Tech Support Command 34-3
Using Traceroute and Other Troubleshooting Tools 34-4
Locating Other Switches 34-5
Configuring an OUI 34-5

CHAPTER 35 **Troubleshooting Fabric Manager Issues** 35-1
Can I Set the Map Layout So It Stays After I Restart Fabric Manager? 35-1
Two Switches Show on my Map, But I Only Have One Switch 35-2
There is a Red Line Through the Switch. What's Wrong? 35-2
There is a Dotted Orange Line Through the Switch. What's Wrong? 35-2
Can I Upgrade Without Losing My Map Settings? 35-2

Are There Any Restrictions When Using Fabric Manager Across FCIP? 35-2
Running Cisco Fabric Manager with Multiple Interfaces 35-3
Specifying an Interface for Fabric Manager Server 35-3
Specifying an Interface for Performance Manager 35-3
Specifying an Interface for Fabric Manager Client or Device Manager 35-4
Configuring a Proxy Server 35-4
Clearing Topology Maps 35-5
Can I Use Fabric Manager in a Mixed Software Environment? 35-5

INDEX

Preface

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

Audience

This guide is for system administrators who intend to use the Cisco Fabric Manager to configure and monitor the switches that build the network fabric.

You should be familiar with the basic concepts and terminology used in internetworking, and understand your network topology and the protocols that the devices in your network can use. You should also have a working knowledge of the operating system on which you are running Fabric Manager, such as Microsoft Windows, Linux, or Solaris.

Organization

This document contains the following chapters:

- **"New and Changed Information"**—Summarizes the new and changed features for the *Cisco MDS 9000 Family Fabric Manager Switch Configuration Guide*.
- **Chapter 1, "Product Overview"**—Presents an overview of the Cisco MDS 9000 Family of multilayer switches and directors.
- **Chapter 2, "Getting Started with Cisco Fabric Manager"**—Provides a brief overview of Fabric Manager components and capabilities, and information on installation and launching the applications.
- **Chapter 3, "Overview of Fabric Manager"**—Provides in-depth descriptions of GUIs and capabilities for Fabric Manager Client, Fabric Manager Server, Device Manager, and Performance Manager.
- **Chapter 4, "Before You Begin"**—Discusses some things to be aware of before installing your switch.
- **Chapter 5, "Obtaining and Installing Licenses"**—Provides information on the Cisco MDS 9000 Family licensing model, license concepts, and license installation and management.
- **Chapter 6, "Initial Configuration"**—Provides initial switch configuration options and switch access information.

- **Chapter 7, "Configuring High Availability"**—Provides details on the high availability feature including switchover mechanisms.
- **Chapter 8, "Software Images"**—Describes how to upgrade Cisco MDS 9000 Family switches, install software image files, use the Flash file system on the supervisor engine, and recover a corrupted bootflash image.
- **Chapter 9, "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 10, "Managing System Hardware"**—Provides details on switch hardware inventory, power usage, power supply, module temperature, fan and clock modules, and environment information.
- **Chapter 11, "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 12, "Configuring Interfaces"**—Explains port and operational state concepts in Cisco MDS 9000 Family switches and provides details on configuring ports and interfaces.
- **Chapter 13, "Configuring Trunking"**—Explains TE ports and trunking concepts.
- **Chapter 14, "Configuring PortChannels"**—Explains PortChannels and load balancing concepts and provides details on configuring PortChannels, adding ports to PortChannels, and deleting ports from PortChannels.
- **Chapter 15, "Configuring and Managing Zones"**—Defines various zoning concepts and provides details on configuring a zone set and zone management features.
- **Chapter 16, "Configuring Inter-VSAN Routing"**—Describes Inter-VSAN Routing.
- **Chapter 17, "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 18, "Configuring Switch Security"**—Discusses the AAA parameters, user profiles, RADIUS authentication, SSH services, and SNMP Security options provided in all switches in the Cisco MDS 9000 Family and provides configuration information for these options.
- **Chapter 19, "Configuring Fabric Security"**—Describes how to configure and manage fabric security.
- **Chapter 20, "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 21, "Configuring Fibre Channel Routing Services and Protocols"**—Provides details and configuration information on Fibre Channel routing services and protocols.
- **Chapter 22, "Configuring IP Services"**—Provides details on IP over Fibre Channel (IPFC) services and provides configuring IPFC, virtual router, and DNS server configuration information.
- **Chapter 23, "Configuring FICON"**—Provides information on configuring and managing FICON with the Cisco MDS 9000 Family.
- **Chapter 24, "Configuring IP Storage"**—Provides details on extending the reach of Fibre Channel SANs by connecting separated SAN islands together via IP networks using FCIP, and allowing IP hosts to access FC storage using the iSCSI protocol.
- **Chapter 25, "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 26, "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 27, "Configuring Traffic Management"**—Provides details on the quality of service (QoS) and Fibre Channel Congestion Control (FCC) features.
- **Chapter 28, "Configuring System Message Logging"**— Describes how to configure system message logging on the Cisco MDS 9000 Family switches.
- **Chapter 29, "Discovering SCSI Targets"**—Describes how the SCSI LUN discovery feature is started and displayed.
- **Chapter 30, "Monitoring Network Traffic Using SPAN"**—Describes the switched port analyzer (SPAN), identifies SPAN sources, specifies filters, explains SPAN Sessions, SD port characteristics, and configuration details.
- **Chapter 31, "Advanced Features and Concepts"**—Describes the advanced configuration features features—time out values, fctrace, fabric analyzer, world wide names, flat FC IDs, loop monitoring, and interoperating switches.
- **Chapter 32, "Configuring Fabric Configuration Servers"**—Describes how the fabric Configuration Server (FCS) feature is configured and displayed.
- **Chapter 33, "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.
- **Chapter 34, "Troubleshooting the Fabric"**—Provides information on using Fabric Manager to troubleshoot your fabric.
- **Chapter 35, "Troubleshooting Fabric Manager Issues"**—Describes some common issues you may experience while using Cisco Fabric Manager, and provides solutions.

Conventions

This document uses the following conventions:

Item	Convention
Commands, keywords, special terminology, and options that should be selected during procedures	**boldface** font
Variables for which you supply values and new or important terminology	*italic* font
Displayed session and system information, paths and file names	`screen` font
Information you enter	**boldface screen** font
Variables you enter	*italic screen* font
Menu items and button names	**boldface** font
Indicates menu items to select, in the order you select them.	**Option > Network Preferences**

Tip Identifies information to help you get the most benefit from your product.

Note Means *reader take note*. Notes identify important information that you should reflect upon before continuing, contain helpful suggestions, or provide references to materials not contained in the document.

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

Warning **Identifies information that you must heed to prevent damaging yourself, the state of software, or equipment. Warnings identify definite security breaches that will result if the information presented is not followed carefully.**

Obtaining Documentation

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

Cisco.com

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

http://www.cisco.com/univercd/home/home.htm

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

Documentation DVD

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

Registered Cisco.com users (Cisco direct customers) can order a Cisco Documentation DVD (product number DOC-DOCDVD=) from the Ordering tool or Cisco Marketplace.

Cisco Ordering tool:

http://www.cisco.com/en/US/partner/ordering/

Cisco Marketplace:

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

Ordering Documentation

You can find instructions for ordering documentation at this URL:

http://www.cisco.com/univercd/cc/td/doc/es_inpck/pdi.htm

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Ordering tool:

  http://www.cisco.com/en/US/partner/ordering/

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 1 800 553-NETS (6387).

Documentation Feedback

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

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

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

We appreciate your comments.

Cisco Product Security Overview

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

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

From this site, you can perform these tasks:

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

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

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

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

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

Reporting Security Problems in Cisco Products

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

- Emergencies—security-alert@cisco.com
- Nonemergencies—psirt@cisco.com

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

Never use a revoked or an expired encryption key. The correct public key to use in your correspondence with PSIRT is the one that has the most recent creation date in this public key server list:

http://pgp.mit.edu:11371/pks/lookup?search=psirt%40cisco.com&op=index&exact=on

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

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

Obtaining Technical Assistance

For all customers, partners, resellers, and distributors who hold valid Cisco service contracts, Cisco Technical Support provides 24-hour-a-day, award-winning technical assistance. The Cisco Technical Support Website on Cisco.com features extensive online support resources. In addition, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not hold a valid Cisco service contract, contact your reseller.

Cisco Technical Support Website

The Cisco Technical Support 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, 365 days a year, at this URL:

http://www.cisco.com/techsupport

Access to all tools on the Cisco Technical Support 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 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 TAC 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 TAC engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

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

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

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

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

Definitions of Service Request Severity

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

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

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

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

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

Obtaining Additional Publications and Information

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

- Cisco Marketplace provides a variety of Cisco books, reference guides, 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

- *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

- 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

New and Changed Information

The table below summarizes the new and changed features for the *Cisco MDS 9000 Family Fabric Manager Switch Configuration Guide*, and tells you where they are documented. If a feature has changed in Release 1.3, a brief description of the change appears in the "Description" column, and that release is shown in the "Changed in Release" column.

Table 1 Documented Features for the Cisco MDS 9000 Family Fabric Manager Switch Configuration Guide

Feature	Description	Changed in Release	Where Documented
All	Updated procedures to remove CLI command references.		
Fabric Manager Switch Tree Change: QoS	Switch->QoS has been moved to Switch->FC->QoS. QoS can only apply to Fibre Channel.	1.3(4)	
Fabric Manager Switch Tree Change: Interfaces	Switch->Interfaces is a new folder which contains: • Port Channels, moved from Switch FC • FC Physical, moved from Switch->FC • FC Logical, moved from Switch->FC • SVC, moved from Switch->FC • Ethernet, gigE and Ether Channels, was previously Switch->IP->Physical Interfaces • Management, mgmt0 and vsan with the ability to create ipfc.	1.3(4)	
Fabric Manager Switch Tree Change: IPFC	Switch->IP->IPFC functionality is now in Switch->Interfaces->Management	1.3(4)	

Table 1 Documented Features for the Cisco MDS 9000 Family Fabric Manager Switch Configuration Guide (continued)

Feature	Description	Changed in Release	Where Documented
Multiple pWWNs to same alias	You can add/associate multiple pWWNs and fWWNs to the same alias name.	1.3(4)	
Quiesce	You can now quiesce/disable a port channel member. This will quiesce a port channel member ISL, and administratively bring down both ports.	1.3(4)	
iSCSI SACK Default	The TCP SACK parameter is enabled by default for iSCSI configurations.	1.3(3)	
Essential Upgrade Prerequisites	Obtaining recommendations based on your current operating environment.	1.3(3)	
iSCSI name restriction	The iSCSI qualified name is restricted to a maximum name length of 223 alphanumeric characters and a minimum length of 16 characters.	1.3(3)	
Rolling upgrades	The Cashing Services Module (CSM) and the IP Storage (IPS) services module use a rolling upgrade install mechanism	1.3(2a)	
Running configuration information	Display Configurations based a specified feature, interface, module, or VSAN.	1.3(1)	
Licensing	Access specified premium features on the switch.	1.3(1)	
Initial Setup Additions	Configure the full zoneset distribution and FC ID persistence features for the entire fabric during initial setup.	1.3(1)	
Automatic image synchronization	The running image is automatically synchronized in the standby supervisor module by the active supervisor module.	1.3(1)	
Standby state	The internal standby state indicates that a switchover is possible when the redundancy state or the supervisor state display standby or HA standby.	1.3(1)	

Table 1 Documented Features for the Cisco MDS 9000 Family Fabric Manager Switch Configuration Guide (continued)

Feature	Description	Changed in Release	Where Documented
Terminal connection options	From the active supervisor module, you can connect to a console terminal, a Telnet terminal, or an SSH terminal.	1.3(1)	
Standby supervisor module boot variables	The software forces the standby supervisor module to run the same version as the active supervisor module.	1.3(1)	
Replacing modules	Ensure that the new module is running the same software version as the rest of the switch. I.	1.3(1)	
Transceiver and calibration information	Display real-time diagnostics information.	1.3(1)	
Buffer-to-Buffer Credit (BB_credit) display	Displays the receive and transmit BB_credit along with other pertinent interface information.	1.3(1)	
PortChannel Quiesce	Use the **quiesce** feature on an ISL to gracefully shutdown an interface without dropping any frames.	1.3(1)	
Zone membership	Assign zone membership criteria is also based on the interface and domain ID, domain ID and port number, and IP address.	1.3(1)	
Inter-VSAN routing (IVR)	Access resources across VSANs without compromising other VSAN benefits.	1.3(1)	
Fabric-Device Management Interface (FDMI)	Enables management of devices using the FDMI feature.	1.3(1)	
AAA server groups	Configure remote AAA servers using server groups.	1.3(1)	
TACACS+ authentication	Use the Terminal Access Controller Access Control System plus (TACACS+) protocol to communicate with remote AAA servers.	1.3(1)	
RADIUS enhancements	Configure multiple RADIUS server groups.	1.3(1)	

Table 1 Documented Features for the Cisco MDS 9000 Family Fabric Manager Switch Configuration Guide (continued)

Feature	Description	Changed in Release	Where Documented
FC-SP DHCHAP	Configure Fibre Channel Security Protocol (FC-SP) authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) provides authentication between Cisco MDS switches and other devices.	1.3(1)	
FI-bre CON-nection (FICON)	Intermix FICON and Fibre Channel Protocol (FCP) traffic on the same switch without compromising scalability, availability, manageability and network security.	1.3(1)	
Fabric Binding	Prevent unauthorized switches from joining the fabric or disrupting current fabric operations.	1.3(1)	
Registered Link Incident Report (RLIR)	Use the RLIR function to send a LIR to a registered Nx-port.	1.3(1)	
Trespass support	Use the trespass feature to enable the export of Logical Units (LUs) from the active to the passive port of a statically imported iSCSI target.	1.3(1)	
Internet Storage Name Service (iSNS)	Use the iSNS services to automate the discovery and management of iSCSI devices.	1.3(1)	
Proxy initiator	Connect all iSCSI initiators through one IPS port to make it appear as one Fibre Channel port per VSAN.	1.3(1)	
FCIP write accelerator	Improve application performance using the FCIP write acceleration feature.	1.3(1)	
FCIP compression	Allow IP packets to be compressed on the FCIP link if this feature is enabled on that link.	1.3(1)	
VSAN membership for iSCSI interfaces	Configure an iSCSI host to be a member of one or more VSANs.	1.3(1)	

Table 1 Documented Features for the Cisco MDS 9000 Family Fabric Manager Switch Configuration Guide (continued)

Feature	Description	Changed in Release	Where Documented
Call Home enhancements	Define a Call Home destination profile, select predefined types of Call Home alerts, or filter messages based on their level of urgency.	1.3(1)	
FC Domain ID changes	Define the default behavior to enable persistent FC IDs globally or for each VSAN.	1.3(1)	
Port rate limiting	Use the port rate limiting feature to control ingress traffic into a Fibre Channel port.	1.3(1)	
Quality of Service (QoS)	Configure four priority levels for service differentiation.	1.3(1)	
Auto-discovery of SCSI targets	Displays automatically discovered SCSI targets.	1.3(1)	
IPS SPAN source	Assign a Switched Port Analyzer (SPAN) source on the IP Storage Services (IPS) module.	1.3(1)	
Per VSAN Time Out Values (TOV)	Configure different TOVs for a specified VSAN with special links like FC or IP tunnels.	1.3(1)	
Deleting directories	Deleting a specified directory deletes the entire directory and all its contents.	All	

Product Overview

The Cisco MDS 9000 Family of multilayer directors and fabric switches offer 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 advanced security and unified management features.

The Cisco MDS 9000 Family provides intelligent networking features such as multiprotocol and multitransport integration, virtual SANs (VSANs), advanced security, sophisticated debug analysis tools, and unified SAN management.

This chapter lists the hardware features for the Cisco MDS 9000 Family and describes its software features.

This chapter contains the following topics:

- Hardware Overview, page 1-1
- Software Features, page 1-4
- Tools for Software Configuration, page 1-10

Hardware Overview

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

- Cisco MDS 9216 multilayer fabric switches contain one fixed integrated supervisor module with 16 Fibre Channel ports and an expansion slot which can support up to 32 additional ports (for a total of 48 ports).
- Cisco MDS 9509 multilayer directors contain two slots for supervisor modules and 7 slots for switching or services modules providing up to 224 ports (32 ports x 7 slots).
- Cisco MDS 9506 multilayer directors contain two slots for supervisor modules and 4 slots for switching or services modules providing up to 128 ports (32 ports x 4 slots).
- Cisco MDS 9140 multilayer switches contains 40 ports (8 full rate ports, 32 host-optimized ports).
- Cisco MDS 9120 multilayer switches contains 20 ports (4 full rate ports, 16 host-optimized ports).

Cisco MDS 9216 Fabric Switch

Cisco MDS 9216 fabric switches share a consistent software architecture with the Cisco MDS 9500 Series in a semi-modular chassis. They consist of the following major hardware components:

- The chassis has two slots, one of which is reserved for the supervisor module. The supervisor module provides supervisor functions and has 16 standard, Fibre Channel ports.
- The backplane has direct plug-in connectivity to one switching module (any type).
- Two redundant, hot-swappable power supplies have AC connections, each of which can supply power to a fully loaded chassis.
- The hot-swappable fan module has four fans managing the airflow and cooling for the entire switch.
- The 1-Gbps or 2-Gbps autosensing Fibre Channel ports support Inter-Switch Links (E ports), extended Inter-Switch Links (TE ports), loop (FL and TL ports), and fabric (F ports) connectivity. Besides Telnet access, a 10/100BASE-T Ethernet port provides switch access and a RS-232 (EIA/TIA-232) serial port allows switch configuration.
- Hot-swappable, small form-factor pluggable (SFP) ports can be configured with either short or long wavelength SFPs for connectivity up to 500m and 10km, respectively.
- The Cisco MDS 9216 supports the IP Storage Services (IPS) module. All IPS modules are configurable for both FCIP and iSCSI operation on a port-by-port basis. Ports configured for FCIP operation can be further configured to support up to three virtual ISL connections.

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

Cisco MDS 9500 Modular Directors

The Cisco MDS 9500 Series includes two multilayer, modular directors:

- The Cisco MDS 9509 Director addresses the stringent requirements of large data center storage environments and consists of the following major hardware components:
  - The chassis has nine slots, two of which are reserved for the supervisor modules.
  - Up to seven hot-pluggable switching or services modules that provide Fibre Channel or Gigabit Ethernet services.
  - The backplane has direct plug-in connectivity to seven switching modules, two integrated supervisor modules, two clock modules, and two power supplies.
  - The hot-swappable fan module has nine fans managing the airflow and cooling for the entire switch.
- The Cisco MDS 9506 Director addresses the stringent requirements of data center storage environments and consists of the following major hardware components:
  - The chassis has six slots, two of which are reserved for the supervisor modules.
  - Up to four hot-pluggable switching or services modules that provide Fibre Channel or Gigabit Ethernet services.
  - The backplane has direct plug-in connectivity to four switching modules, two integrated supervisor modules, two clock modules, and two power supplies.
  - The hot-swappable fan module has six fans managing the airflow and cooling for the entire switch.

These modular directors have the following features:

- Two redundant, hot-swappable power supplies have AC or DC connection, each of which can supply power to the entire chassis.
- Two supervisor modules ensure high availability and traffic load balancing capabilities. Each supervisor module can control the entire switch. The standby supervisor module provides redundancy in case the active supervisor module fails.
- The 1-Gbps or 2-Gbps autosensing Fibre Channel ports support Inter-Switch Links (E ports), Extended Inter-Switch Links (TE ports), loop (FL and TL ports), and fabric (F ports) connectivity. Besides Telnet access, a 10/100BASE-T Ethernet port provides switch access and a RS-232 serial port allows switch configuration.
- Hot-swappable, small form-factor pluggable (SFP) ports can be configured with either short or long wavelength SFPs for connectivity up to 500m and 10km, respectively.
- The Cisco MDS 9500 Series supports the IP Storage Services (IPS) module. All IPS modules are configurable for both FCIP and iSCSI operation on a port-by-port basis. Ports configured for FCIP operation can be further configured to support up to three virtual ISL connections.

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

Cisco MDS 9100 Series Fixed Configuration Fabric Switches

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

- The Cisco MDS 9140 provides 40 ports (8 full rate ports, 32 host-optimized ports)
- The Cisco MDS 9120 is a 20 ports (4 full rate ports, 16 host-optimized ports)

These fixed configuration switches are packaged in a 1 RU enclosures and have the following features:

- Two redundant, hot-swappable power supplies have AC connections, each of which can supply power to the entire chassis.
- Two hot-swappable fan modules with two fans each manage the airflow and cooling for the entire switch.
- The 1-Gbps or 2-Gbps autosensing Fibre Channel ports support Inter-Switch Links (E ports), Extended Inter-Switch Links (TE ports), loop (FL and TL ports), and fabric (F ports) connectivity. Besides Telnet access, a 10/100BASE-T Ethernet port provides switch access.
- Hot-swappable, small form-factor pluggable (SFP) ports can be configured with either short or long wavelength SFPs for connectivity up to 500m and 10km, respectively.

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

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

Software Features

This section provides an overview of the major software features of the Cisco MDS 9000 Family of multilayer directors and fabric switches.

Licensing

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

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. The high availability (HA) software framework includes the following:

- Provides stateful redundancy for supervisor module failure by using dual supervisor modules.
- Ensures nondisruptive software upgrade capability.
- Protects against link failure using the PortChannel (port aggregation) feature. This feature is also available in Cisco MDS 9216 switches and in the Cisco MDS 9100 Series.
- Provides management redundancy using Virtual Router Redundancy Protocol (VRRP). This feature is also available in Cisco MDS 9216 switches and in the Cisco MDS 9100 Series.
- 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 Cisco MDS 9216 switches and in the Cisco MDS 9100 Series.

Switch Reliability

Switches in the Cisco MDS 9000 Family maintain internally controlled reliability services that ensure continued service with no degradation. This reliability service includes the following:

- Provides power-on self testing (POST)
- Detects errors, isolates faults, performs parity checking, and checks illegal addresses
- Enables remote diagnostics using Call Home troubleshooting features
- Displays LEDs that summarize the status of each switching module, supervisor module, power supply, and fan assembly

Virtual SANs

VSANs (virtual SANs) enable higher security and greater scalability in Fibre Channel fabrics. VSANs provide isolation among devices that are physically connected to the same fabric. VSANs allow multiple logical SANs over a common physical infrastructure. VSANs offer the following:

- Traffic isolation—Traffic is contained within VSAN boundaries and devices reside only in one VSAN thus ensuring absolute separation between user groups, if desired.
- Scalability—VSANs are overlaid on top of a single physical SAN. 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 is provided by a configured backup path between the host and the switch.
- Ease of configuration—Devices 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.

Intelligent Zoning

Zoning controls access between devices in a VSAN. Zoning accomplishes the following:

- Partitions devices that use different operating systems. In a heterogeneous environment, it is often necessary to separate servers and storage devices to avoid accidental transfer of information between devices with different operating systems. Such transfers could result in corruption or deletion of data.
- Creates logical subsets of closed user groups. Closed user groups are needed to enforce security or to separate functional areas across the fabric.
- Configures groups of devices that are separate from the rest of the fabric. Based on the assigned zone membership, devices outside the zone cannot access devices internal to the zone.
- Provides temporary access between devices (zone sets). Zone restrictions can be imposed temporarily, and then restored to revert to normal operation, if desired.
- Restricts access to specific logical unit numbers (LUNs) associated with a device.
- Allows members to have only read-only access to the media within a read-only Fibre Channel zone.

Inter-VSAN Routing

Using Inter-VSAN Routing (IVR), resources across VSANs can be accessed without compromising other VSAN benefits. Valuable resources like tape libraries are easily shared across VSANs without compromise. Routes that traverse one or more VSANs across multiple switches can be established, if necessary, to ensure proper interconnections. IVR used in conjunction with FCIP provides more efficient business continuity or disaster recovery solutions.

Trunking

Trunking is the term used to refer to an ISL link that carries one or more VSANs. Trunking ports receive and transmit Extended ISL (EISL) frames. EISL frames carry an EISL header containing VSAN information. Once EISL is enabled on an E port, that port becomes a TE port. The trunking configuration is saved along with the interface information.

PortChannels

PortChannel refers to the aggregation of multiple physical Fibre Channel ports into one logical port to provide high aggregated bandwidth, load balancing, and link redundancy. Up to 16 physical ports can be aggregated into a PortChannel. PortChannels can connect to ports across switching modules. The failure of a port in one switching module does not bring down the logical PortChannel link. Specifically, a PortChannel does the following:

- Increases the aggregate bandwidth on an ISL or EISL 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 a source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) that identify the flow of the frame.
- 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 can contain up to 16 physical links and can span multiple modules for added high availability.

IP Services

Switches in the Cisco MDS 9000 Family support the following IP services:

- IP over Ethernet—These services are limited to management traffic.
- IP over Fibre Channel (IPFC)—IPFC (RFC 2625) specifies how IP packets are transported using encapsulation schemes. By encapsulating IP frames into Fibre Channel frames, management information is exchanged among switches without requiring a separate Ethernet connection to each switch. Each switch includes:
  - Encapsulation for IP and Address Resolution Protocol (ARP) over Fibre Channel.
  - Address resolution uses the ARP server.
- IP routing services—These services include:
  - Ethernet or TCP/IP connection.
  - Static IP routing services to enable management traffic between VSANs.
  - DNS client support.
  - The Network Time Protocol (NTP) server synchronizes the system clocks of network devices.

IP Storage

The Cisco MDS 9000 Family IP services module integrates seamlessly into the Cisco MDS 9000 Family of Multilayer Directors and Fabric Switches. Traffic can be routed between any IP storage port and any other port on a Cisco MDS 9000 Family switch. The Cisco MDS 9000 Family IP Storage Services Module supports the full range of services available on other MDS 9000 Family Switching Modules including VSANs, security, and traffic management. It uses widely known IP to cost-effectively connect to more servers and more locations over greater distances than previously possible. It delivers both Fibre Channel over IP (FCIP) and iSCSI IP storage services and is configurable on a port-by-port basis.

- FCIP highlights
  - Simplifies data protection and business continuance strategies by enabling backup, remote replication, and disaster recovery over WAN distances using open-standard FCIP tunneling.
  - Improves utilization of WAN resources for backup and replication by tunneling up to 3 virtual Inter Switch Links (ISLs) on a single Gigabit Ethernet port.
  - Reduces SAN complexity by eliminating the need to deploy and manage a separate remote connectivity platform.
  - Preserves Cisco MDS 9000 Family enhanced capabilities including VSANs, advanced traffic management, and security across remote connections.
- iSCSI highlights
  - Extends the benefits of Fibre Channel SAN-based storage to IP-enabled servers at a lower cost point than possible using Fibre Channel interconnect alone.
  - Increases storage utilization and availability through consolidation of IP and Fibre Channel block storage.
  - Transparent operation preserves the functionality of legacy storage applications such as zoning tools.
  - Extending the Benefits of Fibre Channel SANs

Call Home

The Call Home feature detects switch failures and sends alerts along with relevant failure information. These alerts are sent through E-mail to a user-specified customer center.

QoS and Congestion Control

Switches in the Cisco MDS 9000 Family provide priority queuing and flow control services.

- The Quality of service (QoS) feature has 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.

- Fibre Channel Congestion Control (FCC)--FCC is a flow control mechanism that alleviates congestion on Fibre Channel networks. Any switch in the network can detect congestion for an output port. The switches sample frames from the congested queue and generate messages about the congestion level upstream toward the source of the congestion. The switch closest to the source, with FCC enabled, can perform one of two actions:
  - Forwards the frames as other vendor switches do.
  - Limits the flow of frames from the port causing the congestion.

SPAN and RSPAN

The switched port analyzer (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.

The Remote SPAN (RSPAN) feature enables you to remotely monitor traffic for one or more SPAN sources distributed in one or more source switches in a Fibre Channel fabric. The SPAN destination (SD) port is used for remote monitoring in a destination switch. A destination switch may be different from the source switch(es) provided that it is attached to the same Fibre Channel fabric. You can replicate and monitor traffic in any remote Cisco MDS 9000 Family switch or director, just as you would monitor traffic in a MDS source switch. This feature is nonintrusive and does not affect network traffic switching for any SPAN source ports.

Switch Management Features

Besides the software features already listed, there are additional management features that fall into the following categories: redundant supervisor module management, fabric management, and security management.

Redundant Supervisor Module Management

Series of multilayer directors support two redundant supervisor modules. They require two supervisor modules to enforce redundant supervisor module management and high availability and restartability.

Table 1-1 Redundant Supervisor Module Management

Product	No. of Supervisor Modules	Slot	Features
Cisco MDS 9216	One module (includes 16 Fibre Channel ports)	Slot 1	2-slot chassis allows one optional switching module in the other slot.
Cisco MDS 9506	Two modules	Slots 5 and 6	6-slot chassis allows any switching module in the other four slots.
Cisco MDS 9509	Two modules	Slots 5 and 6	9-slot chassis allows any switching module in the other seven slots.
Cisco MDS 9120	Not applicable		
Cisco MDS 8149	Not applicable		

When a switch powers up and two supervisor modules are present, the module in slot 5 enters the active mode, while the second module in slot 6 enters the standby mode. All storage management functions occur on the active supervisor module. The standby module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.

See the *Cisco MDS 9500 Series Hardware Installation Guide* for additional information.

Fabric Management

Switches in the Cisco MDS 9000 Family offer fabric management and control through the command-line interface (CLI) by using Telnet, SSH, or a serial console and through the Cisco MDS 9000 Fabric Manager tool by using the Simple Network Management Protocol (SNMP) services:

- SNMP versions 1, 2, and 3 are supported.
- Remote Monitoring (RMON) allows you to specify thresholds and monitor alarms on SNMP variables. Extended RMON alarms are available for supported Management Information Base (MIB) objects. See the *Cisco MDS 9000 Family MIB Reference Guide* for additional information.
- System error message logs (syslogs) are viewed through a console or Telnet session for asynchronous events such as an interface transition. Syslogs are directed to an internal log and optionally to an external server. See the *Cisco MDS 9000 Family System Messages Guide* for additional information.

Security Management

The Cisco MDS 9000 Family of switches offer strict and secure switch management options through switch access security, port security, user authentication, and role-based access.

Switch Access Security

Each switch can be accessed through the CLI or SNMP.

- Secure switch access—Available when you explicitly enable Secure Shell (SSH) access to the switch. SSH access provides additional controlled security by encrypting data, user IDs, and passwords. By default, Telnet access is enabled on each switch.
- SNMP access—SNMPv3 provides built-in security for secure user authentication and data encryption.
- IP Access control lists (IP-ACLs)—Provide basic network security to all switches in the Cisco MDS 9000 Family. IP-ACLs restricts IP-related inband and out-of-band management traffic based on IP addresses (layer 3 and layer 4 information). You can use IP-ACLs to control transmissions on an interface.

Port Security

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 syslog messages.

User Authentication

A strategy known as authentication, authorization, and accounting (AAA) is used to verify identity of, grant access to, and track the actions of remote users. The Remote Access Dial-In User Service (RADIUS) and Terminal Access Controller Access Control System Plus (TACACS+) provide AAA solutions.

Based on the user ID and password combination provided, switches perform local authentication using a local database or remote authentication using AAA server(s). A global, preshared, secret key authenticates communication between the AAA servers. This secret key can be configured for all AAA server groups or for only a specific AAA server. This kind of authentication provides a central configuration management capability.

Role-Based Access

Role-based access control assigns roles or groups (locally through the switch or remotely using AAA servers) to users and limits access to the switch. Access is assigned based on the permission level associated with each user ID. Your administrator can provide complete access to each user or restrict access to specific read and write levels for each command.

From Release 1.2(x), CLI and SNMP in all switches in the Cisco MDS 9000 Family synchronize CLI and SNMP roles. This database contains any role that is created using CLI or SNMP. You can use SNMP to modify a role that was created using CLI and vice versa. Each role in SNMP is the same as a role created or modified through the CLI.

Each role in the role database can be restricted to one or more VSANs as required.

Tools for Software Configuration

You can use one of two configuration management tools to configure your SANs: the CLI and the Cisco MDS 9000 Fabric Manager graphical user interface.

Figure 1-1 Tools for Configuring Software

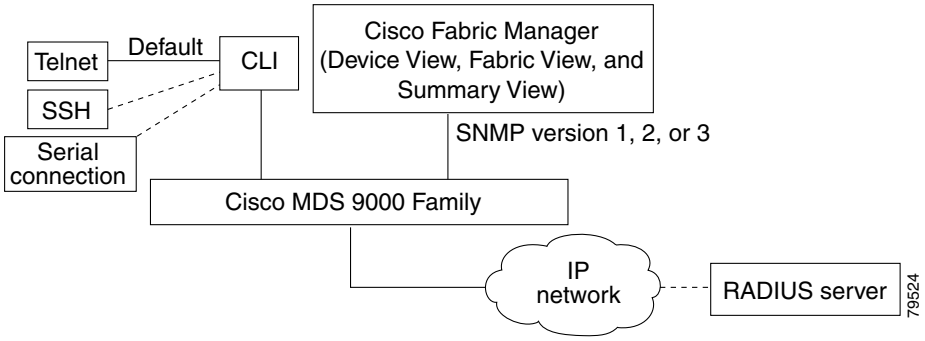

CLI

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.

Cisco MDS 9000 Fabric Manager

The 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 and installed devices. The Cisco Fabric Manager provides three views for managing your network fabric:

- The Device View displays a continuously updated physical picture of device configuration and health conditions for a single switch.
- The Summary View presents real-time performance statistics all active ports and channels on a single switch.
- The Fabric View displays a view of your network fabric, including multiple switches.

The Cisco Fabric Manager provides an alternative to the CLI for most switch configuration commands. The Cisco Fabric Manager is embedded in each switch in the Cisco MDS 9000 Family.

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

Getting Started with Cisco 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. It 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. The Cisco Fabric Manager applications are:

- Fabric Manager Server
- Device Manager
- Fabric Manager Client
- Performance Manager

Fabric Manager Server is the server component of the Cisco Fabric Manager tool set, and must be started before running Fabric Manager. On a Windows PC, Fabric Manager Server is installed as a service. This service can then be administered using the Service Panel in the Control Panel.

The Fabric Manager displays a map of your network fabric, including Cisco MDS 9000 Family switches, third-party switches, hosts, and storage devices.

The Device Manager presents two views of a switch.

- Device View displays a graphic representation of the switch configuration, and provides access to statistics and configuration information for a single switch.
- Summary View displays a summary of xEPorts (Inter-Switch Links), Fx Ports (fabric ports), and Nx Ports (attached hosts and storage) on the switch, as well as FC and IP neighbor devices.

Performance Manager provides detailed traffic analysis by capturing data with SNMP. This data is compiled into various graphs and charts which can be viewed with any web browser.

The Cisco Fabric Manager applications are 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 9000 Family Configuration Guide* or the *Cisco 9000 Family Command Reference*. To learn more about the general capabilities of Cisco Fabric Manager, refer to this document.

This chapter contains the following topics:

- Managing Cisco MDS 9000 Switches, page 2-2
- Storage Management Solutions Architecture, page 2-3
- In-Band Management and Out-of-Band Management, page 2-4
- Installing the Applications, page 2-5
- Launching the Applications, page 2-6
- Using the Management Services Wizard, page 2-7
- A Note on Ports, page 2-7

Managing Cisco MDS 9000 Switches

The Cisco MDS 9000 Family of switches can be accessed and configured in many different ways, and support standard management protocols. The different protocols that are supported in order to access, monitor, and configure the Cisco MDS 9000 Family of switches are described in the following table:

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 70 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	HTTP is only used for the distribution and installation of the Cisco Fabric Manager software. It is not used for communication between the Cisco Fabric Manager and Cisco MDS 9000 Family switches.

Table 2-1 Supported Management Protocols (continued)

Management Protocol	Purpose
ANSI T11 FC-GS3	Fibre Channel-Generic Services (FC-GS)3 in the definition of the management servers defines the Fabric Configuration Server (FCS), which is a standard mechanism to collect information about platforms (end devices) and interconnecting elements (switches) building the fabric. The Cisco MDS 9000 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.
XML/CIM	CIM server support for designing storage area network management applications to run on Cisco MDS SAN-OS.

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 system-oriented view of a 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 interface referred to as 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. IP addresses are resolved to the Fibre Channel address through Address Resolution Protocol (ARP). This feature allows you to build a completely in-band management solution.

Installing the Applications

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

- A supervisor module must be installed on each switch that you want to manage.
- The supervisor module must be configured with the following values using the setup routine or the CLI:
  - IP address assigned to the mgmt0 interface
  - SNMP Credentials (v1/v2 communities, or v3 user name and password), maintaining the same password for all the switches in the fabric. Must be on each PC.

The Cisco Fabric Manager software executables reside on each supervisor module of each Cisco MDS 9000 Family switch 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.

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. When you click the Install buttons on the web page that appears, 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.

New installation options include:

- Upgrade/Downgrade—The installer detects your current version of Fabric Manager and Device Manager, and provides the option to upgrade or downgrade. The default is to upgrade to the latest version of Fabric Manager or Device Manager.
- Autoupgrade—If you always want to run the latest version of Fabric Manager and Device Manager, select “Always autoupgrade, don’t ask me again.” Subsequent upgrades will happen automatically, without prompting.
- Uninstall—Before upgrading or uninstalling Fabric Manager or Device Manager, make sure any instances of these applications have been shut down. Use the Uninstall batch file or shell script to uninstall. Do not delete the .cisco_mds9000 folder as this might make your installation unsafe for upgrades.

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

Step 1 Enter the IP address or host name of the supervisor module 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 web page on Sun Microsystem’s website so you can install it.

The supervisor module HTTP server displays the 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 to reconnect to the supervisor module by reentering the IP address or host name in the Location or Address field of your browser.

Note The recommended version of Java is 1.4.2, although 1.4 and above is supported. To change the JRE version, start Java Web Start and set the Java preferences.

Step 3 Click on any installation link (Fabric Manager or Device Manager).

You see a prompt asking for permission to install the applications on your workstation.

Clicking on a link first runs an installer, which detects the installed version of the software, prompts for upgrades/downgrades and other options if applicable, and runs the application you selected.

All software is installed in a folder called ".cisco_mds9000". On a Windows machine, the pathname is %HOME%\.cisco_mds9000. On a UNIX machine, the pathname is $HOME/.cisco_mds9000.

On a Windows machine, a Cisco MDS program group is created under Start > Programs. This program group contains shortcuts to batch files in the install directory. On a Solaris or Linux machine, shell scripts are created in the install directory.

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

Launching the Applications

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

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

If you double-click on Fabric Manager, the Fabric Manager Server will load (a command line window will appear briefly).

A login screen for Fabric Manager or Device Manager appears.

Step 2 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, accessible through the Device Name(s) drop-down list.

Step 3 Check the **SNMPv3** check box to select SNMP version 3.

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

Step 4 Enter a user name and password.

Step 5 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.

Step 6 Click **Open**.

You see either the Fabric Manager or the Device Manager.

Using the Management Services Wizard

The Management Services Wizard enables you to quickly create and apply initial configuration CLI scripts for the MDS 9000 family. You can run the Management Services Wizard by accessing the switch's supervisor module with a web browser.

To create a configuration script, step sequentially through all the configuration steps by filling in the required fields on each page, and clicking the Next button. You can also selectively configure features by clicking on their category titles in the left frame (e.g., Call Home).

Note When creating configuration scripts this way, you cannot go to any panel with failures in the validation of fields. Use the Enable check box on each panel to enable validation and script generation for that panel. The exception to this is that a Role can only be created if a user is created at the same time.

If you want to change information on a previous page, use the Back button. You can also view and change information you have entered at any time by clicking on the Review category title.

To edit an existing script file, you must open the script file with a text editor, select it, and copy it. In the Management Services Wizard, click on Review under the Script item in the menu tree. Then, paste the script into the text area of the script dialog.

A Note on Ports

For PCs running Fabric Manager Server, Fabric Manager Client, Device Manager, and Performance Manager, certain ports need to be available.

Fabric Manager Client and Device manager use the first available UDP port for receiving SNMP responses. The UDP SNMP Trap local ports are (FM=1162, DM=1163 or 1164). Fabric Manager Client also opens TCP RMI port (9099). If Device Manager is opened from Fabric Manager Client, it listens on the first available UDP port for Fabric Manager requests.

Fabric Manager Server 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 opened, the Fabric Manager client is able to connect to the server. There may be other TCP ports connected to a Fabric Manager client, but they are initiated by server, which is behind the firewall.

Overview of Fabric Manager

This chapter contains descriptions of, and instructions for using, the Cisco MDS 9000 Fabric Manager.

This chapter contains the following topics:

- Launching Cisco Fabric Manager, page 3-1
- Using Fabric Manager, page 3-2
- Discovering and Viewing the Network Fabric, page 3-7
- Controlling Administrator Access with Users and Roles, page 3-7
- Modifying Device Grouping, page 3-7
- Setting Fabric Manager Preferences, page 3-8
- Viewing Reports in Fabric Manager, page 3-9
- Using Device Manager, page 3-10
- Using Performance Manager, page 3-14
- Configuring PM for Use with Cisco Traffic Analyzer, page 3-16
- Stopping Data Collection, page 3-19
- Exporting Data Collection to XML Files, page 3-19
- Removing Data Collection Files from the List, page 3-19

Launching Cisco Fabric Manager

When you click on the Fabric Manager icon, the dialog box allows you to enter the IP address of a computer running the FMServer component. 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 FMServer.

On a Windows PC, you install the FMServer as a service. This service can then be administered using the Service Panel in the Control Panel. The default setting for the FMServer service is that the server is automatically started when the machine is rebooted. You can change this behavior by modifying the properties in the Service panel.

Note If your computer has multiple interface cards (NICs), choose a local interface that can reach Fibre Channel network on clients and on the server.

Using Fabric Manager

The Fabric Manager displays a view of your network fabric, including Cisco MDS 9000 and third-party switches and end devices. To launch the Fabric Manager from your desktop, double-click the Fabric Manager icon and follow the instructions described in the "Launching the Applications" section on page 2-6. The figure below shows the Fabric Manager main window.

Note Changes made using Fabric Manager are applied to the running configuration of the switches you are managing and the changes may not be saved when the switch restarts. After you make a change to the configuration or perform an operation (such as activating zones), Fabric Manager prompts you to save your changes before you exit.

Figure 3-1 Fabric Manager Main Window

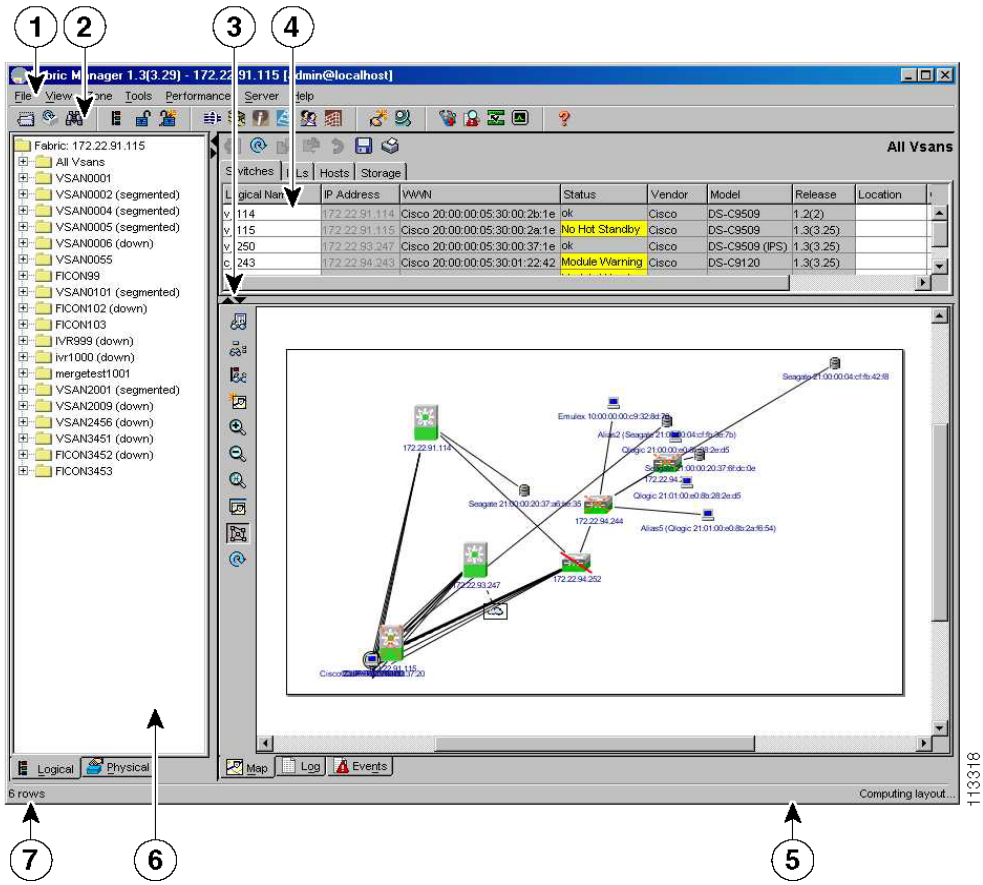

The menu bar at the top of the Fabric Manager window provides access to options, that are organized by menus (see Number 1, Figure 3-1). The toolbar provides icons that duplicate the most commonly used options on the File, Tools, and Help menus (see Number 2, Figure 3-1).

The main window has a menu bar, toolbar, message bar, status bar, and three panes:

- Logical/Physical pane—Displays a tree of configured VSANs and zones on the Logical tab and a menu tree of available configuration tasks on the Physical tab (see Number 6, Figure 3-1).
- Information pane—Displays information about whatever option is selected in the menu tree (see Number 4, Figure 3-1).
- Map pane—Displays a map of the network fabric, including switches, hosts, and storage. It also provides tabs for displaying log and event data (see Number 3, Figure 3-1).

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

Menu Bar, Toolbars, and Status 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 Map pane. The menu bar provides the following menus:

- File—Open a new fabric, rediscover the current fabric, locate switches, set preferences, print the map, and export the Map pane log.
- View—Change the appearance of the map (these options are duplicated on the Map pane toolbar).
- Zone—Manage zones, zone sets, and Inter-VSAN Routing (IVR).
- Tools—Verify and troubleshoot connectivity and configuration, as described in the “Analyzing Switch Fabric Configuration” section.
- Performance—Run and configure Performance Manager and Cisco Traffic Analyzer, and generate reports.
- Server—Run administrative tasks on clients and fabrics.
- Help—Display online help topics for specific dialog boxes in the Information pane.

The Fabric Manager main toolbar provides buttons for accessing the most commonly used menu bar options. The Map pane toolbar provides buttons for managing the appearance of the map. The Information pane toolbar provides buttons for editing and managing the Information pane.

The status bar shows the last entry displayed by the discovery process, and the possible error message on the right side. It displays a dialog 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.

Logical/Physical Pane

Use the Logical tab on the Logical/Physical pane to manage virtual SAN attributes (e.g., zones) in the currently discovered fabric.

To manage zones, right-click one of the folders in the VSAN tree and click Edit Local Zone Database from the pop-up menu. You see the Edit Local Zone Database dialog box.

Use the Physical tab on the Logical/Physical pane to display a menu tree of the options available for managing the switches in the currently discovered fabric.

To select an option, click a folder to display the options available and then click the option. You see the dialog box for the selected option in the Information pane.

The Physical tree provides the following main folders:

- Switches—View and configure hardware, system, licensing, and configuration files.
- Interfaces—View and configure FC Physical, FC Logical, Ethernet, SVC, and Port Channels interfaces.
- FC—View and configure Fibre Channel network configurations.
- IP—View and configure IP storage and IP services.
- Events—View and configure events, alarms, thresholds, notifications, and informs.
- Security—View and configure MDS management and FC-SP security.
- Connectivity—View and configure ISLs, Hosts, and Storage components.

Information Pane

The Information pane displays tables or other information associated with the option selected from the menu tree. The Information pane toolbar provides buttons for performing one or more of the following operations:

- Apply Change— Apply configuration changes.
- Refresh Value— Refresh table values.
- Copy (Ctrl-C)—Copy data from one row to another.
- Paste (Ctrl-V)—Paste the data from one row to another.
- Undo Changes (Ctrl-Z) —Undo the most recent change.
- Export—Export and save information to a file.
- Print Table—Print the contents of the Information pane.

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 you select. They are activated or deactivated (grayed) according to the field or other object that you select in the Information pane.

Map Pane

The Map pane shows the graphical representation of your fabric.Table 3-1 explains the graphics you may see displayed, depending on which devices you have in your fabric.

Table 3-1 Fabric Manager Graphics

Icon or Graphic	Description
	Director Class MDS 9000
	Non-director Class MDS 9000
	Generic FC Switch
	Cisco SN5428
	A line through a device indicates that the device is not manageable
	An “X” through a device or link indicates that the device is down or that the connection is down
	FC HBA (or enclosure)
	FC Target (or enclosure)
	iSCSI Host
	Fibre Channel ISL and Edge connection
	Fibre Channel Port Channel
	IP ISL and Edge connection
	IP Port Channel
	FC Loop (Storage)

Table 3-1 Fabric Manager Graphics (continued)

Icon or Graphic	Description
	IP Cloud (Hosts)
	Any device, cloud, or loop with a box around it means that there are hidden links attached

There are three tabs on the bottom of the Map pane:

- Map—Displays a graphical view of the network fabric with switches, hosts, and storage subsystems.
- Log—Displays messages that describe Fabric Manager operations, such as fabric discovery.
- Events—Displays information about the SNMP traps received by the management station.

When viewing large fabrics in the Map pane, it is helpful to:

- 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

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

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 Map pane toolbar or choose Clear Highlight from the pop-up menu.

Discovering and Viewing the Network Fabric

Cisco Fabric Manager collects information on the fabric topology through SNMP queries to the switches connected to Fabric Manager. The switch replies after having discovered all devices connected to the fabric by using the information coming from its FSPF technology database and the Name Server database, and collected using the Fabric Configuration Server's request/response mechanisms defined by the FC-GS3/4 standard. When you start the Fabric Manager, you enter the IP address (or host name) of a "seed" switch for discovery.

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

Controlling Administrator Access with Users and Roles

Cisco MDS 9000 Family switches support role-based management access whether using the CLI or the 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 new 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.

Modifying Device Grouping

Because of not all the devices are capable of responding to FC-GS3 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 in order to have them represented by a single icon on the map, follow these steps:

Step 1 Select **Storage** or **Hosts** from the Fabric Manager's Physical tree in the Navigation pane.

The end devices appear in the Information pane.

Step 2 Click on the Name field for one of the devices you want to be in the enclosure.

Step 3 Enter a name to identify the new enclosure's icon on the Fabric Manager Map pane.

Step 4 Click once on the Name field for that device. To select more than one Name, hold down 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 name into the Name field for that device.

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

Setting Fabric Manager Preferences

To set your preferences for the behavior of the Fabric Manager, choose File > Preferences from the Fabric Manager menu bar. The Preferences dialog box appears.

This dialog box has the following tabs, which let you set your preferences for different components of the application:

- General
- Discovery
- Map

The default General preferences for Fabric Manager are:

- **Show Switch Name by IP**—This displays the IP addresses of the switches in the Map pane, rather than the DNS or Logical Name.
- **Show WWN Vendor by Prepend Name**—The other options are to show the vendor by replacing vendor bytes, or not to show it at all.
- **Append Enclosures to End Device Names**—The default setting for this value is OFF.
- **Show Shortened iSCSI Names**—The default setting for this value is OFF.
- **Show Timestamps as Date/Time**—The default setting for this value is ON.
- **Use Secure Shell instead of Telnet**—The default setting for this value is OFF. When set to ON, you must enter a path for your secure shell client.
- **Confirm Deletions**—The default setting for this value is ON.
- **Export Tables with Tab-Delimited Format**—The other option is to export with XML format.

The default SNMP preferences for Fabric Manager are:

- **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.
- **Discover LUN by Host OS**—The default setting for this value is All. The other options are Windows, AIX, Solaris, Linux, and HPUX

The default Map preferences for Fabric Manager are:

- **Display Unselected VSAN Members**—This displays the unselected VSAN members in the Map pane. The default setting for this value is ON.
- **Display End Devices**—This displays the fabric's end devices in the Map pane. The default setting for this value is ON.
- **Display End Device Labels**—This displays the fabric's end device labels in the Map pane. The default setting for this value is ON.
- **Expand Loops**—This displays the loops in the fabric as individual connections in the Map pane. The default setting for this value is OFF.
- **Expand Multiple Links**—This displays multiple links in the Map pane as separate lines rather than as one thick line. The default setting for this value is ON.
- **Open New Device Manager Each Time**—This opens a new instance of Device Manager each time you invoke it from a switch in your fabric. The default value is OFF, which means only one instance of Device Manager will be open at a time.

- **Layout New Devices Automatically**—This automatically places new devices in the Map 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 End Devices**—The default setting for this value is 30. You can enter any number in this field. Enter 0 to disable Quick Layout.
- **Override Preferences for Non-default Layout**—The default setting for this value is ON.
- **Automatically Save Layout**—When this option is enabled, any changes in layout are automatically saved. The default setting for this value is ON.

Viewing Reports in Fabric Manager

The Fabric Manager provides a series of reports, showing various information in tabular form. When you select one of these options, you see the available information in tabular form in the Information pane of the Fabric Manager main window. The table below describes the reports provided by each option.

Table 3-2 Fabric Manager Reports

Reports	Description
ISL Statistics	Click on **Connectivity > ISLs > Statistics** in the Physical tab of the Fabric Manager Logical/Physical pane to display information about the Inter-Switch Links in the currently discovered fabric. You can use the controls at the top of the table to change the Poll Interval and Scale parameters.
ISL	Choose **Connectivity > ISLs** in the Physical tab of the Fabric Manager Logical/Physical pane to display information about the Inter-Switch links in the currently discovered fabric.
Switches	Choose **Switches** in the Physical tab of the Fabric Manager Logical/Physical pane to display information about the switches in the currently discovered fabric.
Hosts	Choose **Connectivity > Hosts** in the Physical tab of the Fabric Manager Logical/Physical pane to display information about the hosts in the currently discovered fabric.
Storage	Choose **Connectivity > Storage** in the Physical tab of the Fabric Manager Logical/Physical pane to display information about the links to hosts and storage in the currently discovered fabric.
LUNs	Choose **Connectivity > Storage > LUNs** in the Physical tab of the Fabric Manager Logical/Physical pane to display information about the LUNs in the currently discovered fabric.

Using Device Manager

Device Manager provides a physical representation of your switch chassis, with the modules, ports, power supplies, and fan assemblies. The menu bar at the top of the Device Manager window provides access to options, organized into menus that correspond to the menu tree in Fabric Manager.

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

- Green—The port is up.
- Brown—The port is administratively down.
- Red—The port is down or has failed.
- Gray—The port is unreachable.

Launching Device Manager from Fabric Manager

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

To launch the Device Manager from your desktop, double-click the Device Manager icon and follow the instructions described in the “Launching the Applications” section.

To launch Device Manager from Fabric Manager, right-click the switch you want to manage on the Fabric Manager map and click Device Manager from the pop-up menu that appears. The Device Manager main window is shown below.

Device Manager can also be started by double-clicking on a switch in the Fabric Manager topology view, or by selecting a switch in the Fabric Manager Map page and choosing the Device Manager option from the Tools menu.

Figure 3-2 Device Manager, Device Tab

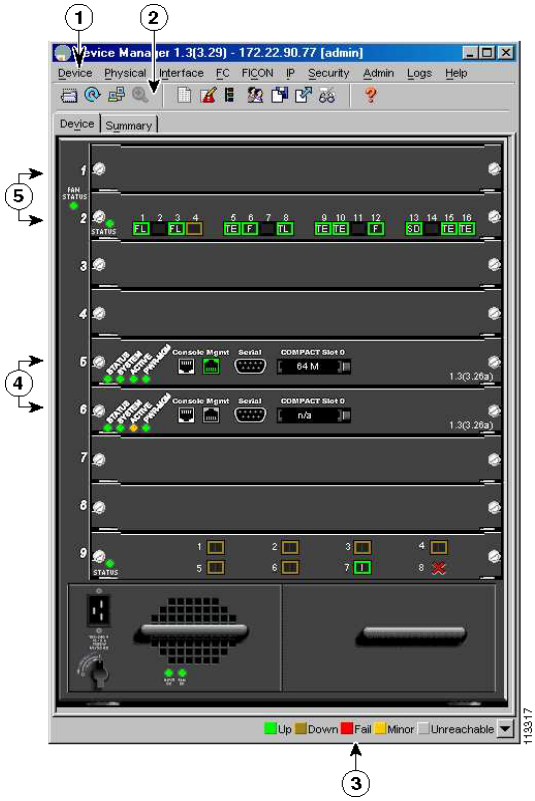

Using Summary View

Click the Summary tab on the Device Manager main window to see a summary of enabled interfaces on a single switch, as well as FC and IP neighbor devices. All logical interfaces are shown in a dropdown list at the top of the Summary view.

The Summary View displays attributes for a single switch, such as port speed, link utilization, and other traffic statistics. It has the same menu bar and toolbar buttons as the Device View.

To monitor traffic for selected objects, click the Monitor icon. To display detailed statistics for selected objects, click the Detailed Statistics icon.

The Summary View provides the same menus and options that are available from the Device View.

Note You can access the field descriptions for the windows or dialog boxes in this procedure in the Reference section of the Device Manager help system.

Comparing Device Manager to Fabric Manager

The menu bar at the top of the Device Manager contains the same menus as the Fabric Manager menu tree.

For information about the options provided by these menus, see the "Logical/Physical Pane" section on page 3-3. The Device menu provides the following options:

- Open—Open the Device Manager for a different switch.
- Open Last—Open the Device Manager for the most recently managed switch.
- Preferences—Set management preferences for controlling the behavior and appearance of the Device Manager.
- Refresh—Update the current display.
- Command Line Interface—Open a Telnet/SSH session with the current switch.
- Exit—Close the Device Manager application.

The tables in the Fabric Manager roughly correspond to the dialog boxes that appear in Device Manager. However, the Fabric Manager tables show values for multiple switches and so the first column identifies the specific switch. The Device Manager dialog box shows values for a single switch, while the Fabric Manager shows the same values for one or more switches.

The toolbar on the Device Manager dialog box provides the same options as the toolbar on the Information pane in Fabric Manager, as summarized here:

- Create—Insert a new row into a table (if applicable).
- Delete Row—Delete the selected row from a table (if applicable).
- Copy (Ctrl-C)—Copy data from one row to another.
- Paste (Ctrl-V)—Paste the data from one row to another.
- Apply Changes—Apply configuration changes. (Note: After making changes you must save the configuration. Otherwise, the changes will be lost when the device is restarted.)
- Refresh Values—Refresh table values.
- Reset Changes (Ctrl-Z)—Undo the most recent change.
- Print table—Print the contents of the Information pane.

Tip You can copy values from one cell in a table to the rest of the column. Copy the value to the clipboard, hold down the shift key while pressing the down arrow key (or click on the bottom cell in the column). Then paste the value to all the selected cells and click Apply.

When you click the Create button, you see a dialog box that lets you enter the values required for the specific table. As you can see the fields and options are the same from both views, but the appearance of the window may vary slightly. For instance, the dialog box from Fabric Manager may have an option for selecting a specific switch, while the dialog box from Device Manager may have additional port-level detail.

Performing Device Management

Most tasks that you can perform with Device Manager can also be performed for multiple switches using the Fabric Manager. However, Device Manager may be more convenient to use when you are working with a single switch. Also, the Device Manager provides more detailed information for verifying or troubleshooting device-specific configuration than what is available from the Fabric Manager.

The Device View provides a graphic representation of a Cisco MDS 9000 switch, including the installed switching modules, services modules, supervisor modules, and the status of each port within each module. You can use the Device View to perform any switch-level configuration tasks including the following:

- Manage ports, Port Channels, and trunking
- Manage SNMPv3 security access to switches
- Manage CLI security access to switches
- Manage alarms, events, and notifications
- Save and copy configuration files and software images
- View hardware configuration
- View chassis, module, and port status and statistics

Summary View provides a way of monitoring all of the ports on the switch, categorized by operative modes (Fx-Ports and E-Ports).

When you click the Summary tab on the Device Manager window, you see the Summary View, which provides summary information about the interfaces on a single switch.

Managing Ports

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 on 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**. On the dialog box that appears, in the Trunk column, right-click the current value 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 the"About PortChanneling and Trunking" section on page 14-3. 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.

Setting Device Manager Preferences

From the Device menu, choose **Preferences** to set your preferences for the behavior of the Device Manager application.

Note You can access the field descriptions for the windows or dialog boxes in this procedure in the Reference section of the Device Manager help system.

Using Performance Manager

Performance Manager monitors network device statistics historically, and provide 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.

Performance Manager Architecture

The Performance Manager has three parts:

- Definition—Use a configuration wizard to create a configuration file
- Collection—Performance Manager reads the configuration file and collects the desired information
- Presentation—Performance Manager generates web pages to present the collected data Performance Manager can collect a variety of data, about these fabric components: ISLs, host ports, storage ports, route flows, and site-specific statistical collection areas.

Creating a PM Configuration File

Performance Manager has a Configuration File Wizard, which steps you through the process of creating configuration files.

To create a configuration file, perform the following steps:

Step 1 Choose **Performance > Create Collection** in Fabric Manage to launch the wizard.

Step 2 Select the VSANs from which you want to collect data.

Step 3 Check the types of SAN objects for which you want to collect data.

Step 4 If you want to ignore flows with Zero counter values, check that check box.

Step 5 If you are using Cisco Traffic Analyzer, enter the URL where it is located on your network.

Step 6 Click **Next** to review the collected data.

Step 7 Enter the name of the file (the default is the switch's IP address with a .XML suffix).

Step 8 Select the definitions that you wish to remove, then click **Finish** to create the configuration file.

Note The FV Interface mentioned at the bottom of the second screen of this wizard is the Veritas virtualization interface. It is documented in the Veritas documentation.

Collecting the Data

One year's worth of data for two variables (Rx and Tx Bytes) requires an rrd file size of 76K. 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 + 50 days)
- 300 samples of 1 day (above + 300 days, rounded up to 365)

A 1000-port SAN requires 76MB for a year's worth of historical data. If there were 20 switches in this SAN with equal distribution of fabric ports, about 2-3 SNMP packets per switch would be sent every 5 minutes for a total of about 100 total request/response SNMP packets required to monitor the data.

Flows, because of their variable counter requests, are more difficult to predict. But as a rule of thumb, each extra variable adds another 38K.

The Performance Manager collector is designed to run as a background process on the various supported OSs. On Microsoft Windows, it runs as a service.

Presenting the Collected Data

The Summary page presents the top 10 Hosts, ISLs, Storage, and Flows by average throughput for the last 24 hour period. This period changes on every polling interval ñ this is unlikely to change the average significantly, but it could affect the maximum value. The intention is to provide a quick summary of the fabric's bandwidth consumption and highlight any hotspots.

- Clicking on any Host, Storage, ISL, or Flow title will provide a view of traffic over the past day for all Hosts, Storage, ISLs, or Flows respectively.
- Clicking on a host port from the summary page will provide you with a similar detail page. If flows exist for that port, you could see which storage ports it was sending data to.
- Clicking on the ISLs link from the summary page will list the daily traffic charts for all monitored ISLs in the fabric.

Exporting and Importing Data

You can export an rrd file to XML with the command:

```
pm xport <rrdFile> <xmlFile>
```

This will produce an XML format that rrdtool is capable of reading with the command:

```
rrdtool restore filename.xml filename.rrd
```

You can import an XML with the command:

```
pm restore <xmlFile> <rrdFile>
```

This will read the XML export format that rrdtool is capable of writing with the command:

```
rrdtool xport filename.xml filename.rrd
```

Integration with Cisco Traffic Analyzer

SNMP and Performance Manager can only provide a top-level view of what data the fabric is carrying. The Cisco MDS 9000 switch has no LUN-level flow counters, and cannot count SCSI commands. In order to view this detailed information, it is necessary to look at the data on a SPAN destination port with the help of the Cisco Traffic Analyzer, which uses the Cisco Port Adapter Analyzer.

Cisco Traffic Analyzer must be downloaded and installed separately.

Caution The Cisco Traffic Analyzer for Fibre Channel throughput values are not accurate when used with the original Cisco Port Adapter Analyzer if data truncation is enabled. The A version of the Cisco Port Adapter Analyzer is required to achieve accurate results with truncation, because it adds a count that enables the Cisco Traffic Analyzer to determine how many data bytes were actually transferred.

Configuring PM for Use with Cisco Traffic Analyzer

Performance Manager works in conjunction with the Cisco Traffic Analyzer to allow you to monitor and manager the traffic on your fabric. The figure below is a graphic representation of how Performance Manager works with the Cisco Traffic Analyzer to monitor traffic on your fabric.

Figure 3-3 Overview of Performance Manager/Cisco Traffic Analyzer Configuration

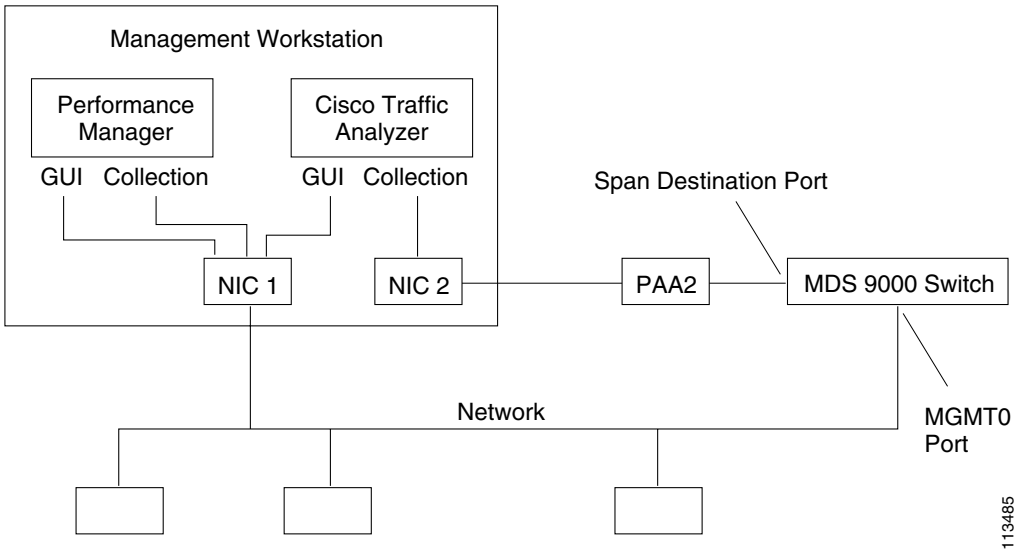

Note See Chapter 30, "Monitoring Network Traffic Using SPAN" for information on configuring the settings for your span destination port. It is important that the data you collect through this port matches the data collected by Performance Manager through the MGMT0 port. If the data does not match, you will not be able to view the Cisco Traffic Analyzer information through a Traffic Analyzer link on the detail page of a Performance Manager report.

For information on setting up the Cisco Traffic Analyzer, refer to the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note.*

To configure Performance Manager to work with the Cisco Traffic Analyzer, perform 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 You will need 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 3 Start the Cisco Traffic Analyzer.

1. From the Fabric Manager Performance menu, choose **Traffic Analyzer > Open**.
2. 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).

3. Click **OK**.
4. From the Fabric Manager Performance menu, choose **Traffic Analyzer > Start**.
5. 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.
6. Click **OK**.

Step 4 Create the flows you want Performance Manager to monitor, using the Flow Configuration Wizard.

Step 5 Define the data collection you want Performance Manager to gather, using the Performance Manager Configuration Wizard.

1. Select the VSAN you want to collect information for, or select All VSANs.
2. Check the types of items you want to collect information for (hosts, ISLs, Storage devices, and Flows).
3. 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.
4. Click **Next**.
5. Review the Data Collection to make sure this is the data you want to collect.
6. 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 in order for your changes to take effect.

Step 6 To generate a report, choose **Performance > Reports**.

You see a list of XML files. These files are the data collection files you specified in the Performance Manager Configuration Wizard.

Note It takes about five minutes to collect enough data to generate a report. Do not attempt to generate a report in Performance Manager during the first five minutes of collection.

Step 7 Choose a file for which you want to generate a report.

In about five minutes, an HTML report appears in your default web browser.

Step 8 To view the Cisco Traffic Analyzer information, click the **Cisco Traffic Analyzer** link at the top of the Host or Storage detail pages.

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 "Using Performance Manager" section on page 3-14.

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*.

Stopping Data Collection

You can stop a data collection process in Windows using the services panel. Right click on the Cisco Performance Manager service and select Stop.

On a Unix machine, enter the following command:

```
$HOME/.ciscomds9000/bin/pm.sh stop
```

Exporting Data Collection to XML Files

To export the collection to an XML file, enter the following command at the operating system command line prompt:

```
$HOME/.ciscomds9000/bin/pm.bat xport xxx yyy
```

Removing Data Collection Files from the List

To remove a data collection file from the list, edit the pm.txt file and comment out the line (#), or remove the line entirely.

Before You Begin

This chapter lists the information you need to have before you begin using your MDS 9000 Switch. For information on setting up the switch and doing an initial configuration, refer to the *Cisco MDS 9000 Family Configuration Guide*.

This chapter contains the following topics:

- About Flash Devices, page 4-1
- Switch Roles, page 4-2
- Using Valid Formats and Ranges, page 4-2

About Flash Devices

Every switch in the Cisco MDS 9000 Family contains one internal bootflash. The Cisco MDS 9500 Series additionally contains one external CompactFlash called slot0. (See Figure 4-1 and Figure 4-2.)

Figure 4-1 Flash Devices in the Cisco MDS 9000 Supervisor Module

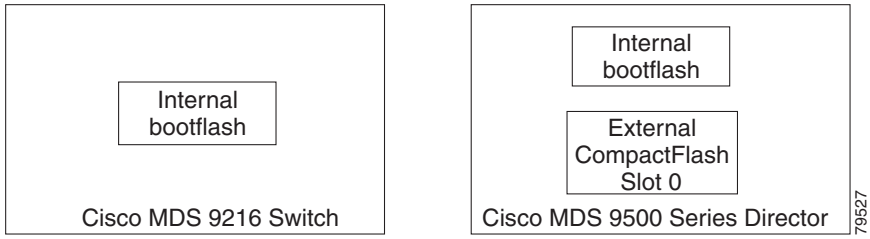

Figure 4-2 External CompactFlash in the Cisco MDS 9000 Supervisor Module

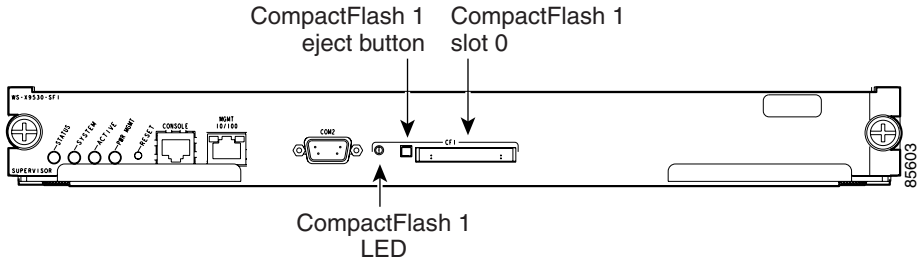

Internal bootflash:

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 directories within the internal bootflash: file system.

- The volatile: directory which provides temporary storage, and is also the default. Files in temporary storage (volatile:) are erased when the switch reboots.
- The bootflash (nonvolatile storage): directory which provides permanent storage. The files in bootflash are preserved through reboots and power outages.

External CompactFlash (Slot0)

Cisco MDS 9500 Series directors contain an additional external CompactFlash called slot0:

The external CompactFlash, an optional device for MDS 9500 Series directors, can be used for storing software images, logs, and core dumps.

Switch Roles

By default, two roles exist in all switches:

- Network operator—Has permission to view the configuration.
- Network administrator—Has permission to execute all commands and to set up to 64 permission levels based on user roles and groups.

When you execute a command, perform command completion, or obtain context sensitive help, the switch software allows the operation to progress if you have the correct permission as specified in the description of the command.

Using Valid Formats and Ranges

Note Do not enter ellipsis (...), vertical bars (|), less than or greater than (< >), brackets ([ ]), or braces ({ }) in any formats or ranges. These characters have special meaning in SAN-OS text strings.

Table 4-1 Valid Formats and Ranges

Address	Description	Valid Format Example	Range
MAC address	6 bytes in hexadecimal format separated by colons (not case-sensitive)	00:00:0c:24:d2:Fe	--

Table 4-1 Valid Formats and Ranges (continued)

IP address	32 bytes, written as 4 octets separated by periods (dotted decimal format) that are made up of a network section, an optional netmask section, and a host section.	126.2.54.1	--
VSAN	Integer that specifies the VSAN.	7	1 to 4093
VLAN	Integer that specifies the VLAN	11	1 to 4093
Port WWN (p WWN)	Eight hexadecimal numbers separated by colons (not case-sensitive).	12:34:56:78:9A:BC:dE:F1	--
Node WWN (n WWN)	Eight hexadecimal numbers separated by colons (not case-sensitive).	12:34:56:78:9A:BC:dE:F1	--
LUN	8 bytes in hexadecimal format separated by colons. A minimum of two hex characters are acceptable. The valid format is hhhh[:hhhh[:hhhh[:hhhh]]]	64 (100d = 64h)	--
F C ID	Six character hexadecimal value prepended by 0x.	0xabc123	--
Domain ID	Integer that specifies the domain.	7	1 to 239
Timers	Integer that specifies timers in milliseconds for latency, FC time out values (TOV).	100	0 to 2147483647
Switching module	Slot in which the applicable switching module resides.	1	1 to 15
Switch priority	Integer specifying switch priority.	5	1 to 254
Channel group	Integer that specifies a PortChannel group addition.	1	1 to 100

Table 4-1 Valid Formats and Ranges (continued)

Fabric Shortest Path First (FSPF)	Integer that specifies the hold time (in milliseconds) before making FSPF computations.	1000	0 to 65535
Fabric Analyzer	The allowed range for the frame size limit in bytes.	64	64 to 65536
Fabric Analyzer captures	An example of 10 frames, limits the number of frames captured to 10.	10	0 to 2147483647
FCIP profile	Integer that specifies the FCIP profile	101	1 to 255
TCP retransmit time	Integer that specifies the minimum retransmit time for the TCP connection in milliseconds	300	250 to 5000
Keepalive timeout	Integer that specifies the TCP connection's keepalive timeout in seconds.	60	1 to 7200
TCP retransmissions	Integer that specifies the maximum number of TCP transmissions.	6	1 to 8
PMTU	Integer that specifies the path MTU reset time in seconds	90	60 to 3600
TCP buffer size	Integer that specifies the advertised TCP buffer size in KB.	5000	0 to 8192
Traffic burst size	Integer that specifies the maximum burst size in KB.	30	10 to 100
Peer TCP port	Integer that specifies the TCP port number	3000	0 to 65535
Acceptable time difference	Integer that specifies the acceptable time difference in milliseconds for a packet being accepted.	4000	1 to 60,000

Table 4-1 Valid Formats and Ranges (continued)

iSCSI pWWN allocation	Integer that specifies the number of pWWNs that must be allocated to an iSCSI initiator.	2	1 to 64
CDP refresh and hold time	Integer that specifies the refresh time interval and the hold time in seconds for the CDP protocol.	60	5 to 255

Obtaining and Installing Licenses

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

This section contains information related to licensing types, licensing procedure, license installation, and license management for the Cisco MDS SAN-OS software.

This chapter contains the following topics:

- License Terminology, page 5-1
- Licensing Model, page 5-2
- Licensing High Availability, page 5-4
- Options to Install a License, page 5-4
- Obtaining a Factory-Installed License, page 5-4
- Performing a Manual Installation, page 5-5
- Obtaining License Key Files, page 5-5
- Installing Licenses, page 5-6
- Viewing License Information in Fabric Manager, page 5-8
- Viewing License Information in Device Manager, page 5-9
- Removing Licenses, page 5-9
- Updating Licenses, page 5-10
- License Expiry Alerts, page 5-10
- Moving Licenses Between Switches, page 5-11

License Terminology

The terms used in this chapter are explained in this section.

- 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.
- License expiry—The time span during which a licensed feature is valid. The software tracks all licenses and sends periodic alerts before shutting down the licensed feature.

a. Counted license—The number of usage instances for a licensed feature.

- Licensed application—A software feature that requires a license in order to be used.
- License enforcement—Mechanism to prevent features 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 ID—A host ID is a unique chassis serial number that is specific to each Cisco MDS switch.
- Proof of purchase—Also known as the Claim Certificate. A document entitling it's rightful owner to use licensed feature(s) of the Cisco MDS 9000 Family of Multilayer Directors and Fabric Switches on one MDS switch as described in that document.
- Product Authorization Key (PAK)—Using the PAK, you can 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 via 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 licenses—Counted licenses refer to 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.
- Incremental licenses—Incremental refers to adding other licensed features 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—An evaluation license is 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).
- Permanent license—A license that is not time bound (does not have an expiry date) is called a permanent license.
- Grace period—The amount of time an application can continue functioning without a license. The grace period is set to 120 days from the first occurrence of using any licensed feature without a license. The grace period starts with the first checkout and will be counted only for the days when that feature is used. If you do not use this feature, the grace period stops incrementing.
- Support—If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco, contact Cisco Technical Support at this URL:

Licensing Model

The licensing model defined for the Cisco MDS product line has two options:

- Feature-based licensing—Features that are applicable to the entire switch. The cost varies based on a per-switch usage. lists the feature-based license packages.
- Module-based licensing—Features that require additional hardware modules. The cost varies based on a per-module usage. Some examples are the IPS module using the FCIP feature or a Fibre Channel switching module using the FICON feature.

Table 5-1 lists the licenses and their features.

Table 5-1 Feature-Based Licenses

Feature License	Features
Standard package (free--no license required)	• FCP, SSH, SFTP, and iSCSI protocols • Fabric Manager and Remote monitoring (RMON) • VSANs, High availability, PortChannel, and Zoning • Fibre Channel Congestion Control (FCC) • Virtual Output Queuing (VOQ) • Diagnostics (SPAN, RSPAN, and FC Analyzer) • SNMP v3, Role-based access control, RADIUS • Call Home and Interoperability modes • IP access control lists (ACLs) • Terminal Access Controller Access Control System (TACACS+) • Fabric-Device Management Interface (FDMI) • Internet Storage Name Service (iSNS) client.
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) • Enhanced VSAN routing--inter-VSAN routing
SAN extension over IP (SAN_EXTN_OVER_IP)	• FCIP protocol • FCIP compression • FCIP write acceleration

Table 5-1 Feature-Based Licenses (continued)

Mainframe (MAINFRAME_PKG)	• FICON protocol and CUP management • FICON VSAN and intermixing • Switch cascading • Fabric Binding
Fabric Manager Server (FM_SERVER_PKG)	• Multiple physical fabric management • Centralized fabric discovery services • Continuous MDS health and event monitoring • Long term historical Fibre Channel Performance monitoring • Performance reports and charting for hotspot analysis

Licensing High Availability

Like any other Cisco MDS SAN-OS feature, 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.
- When a licensed feature is enabled without a license key, the MDS switch enables the feature and starts a counter on the grace period. You then have 60 days to install the appropriate license keys or disable the use of that feature. If at the end of the 60 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, you have two options:

- To have the licenses preinstalled in the factory.
- To install the licenses yourself by following the manual process.

If you already have an existing switch, follow the manual process.

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, contact Cisco Technical Support at this URL:

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 for future use.

Step 3 Start using the switch and the installed licenses 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. Figure 5-1 maps out ways to obtain license key files.

Figure 5-1 Obtaining a License Key File

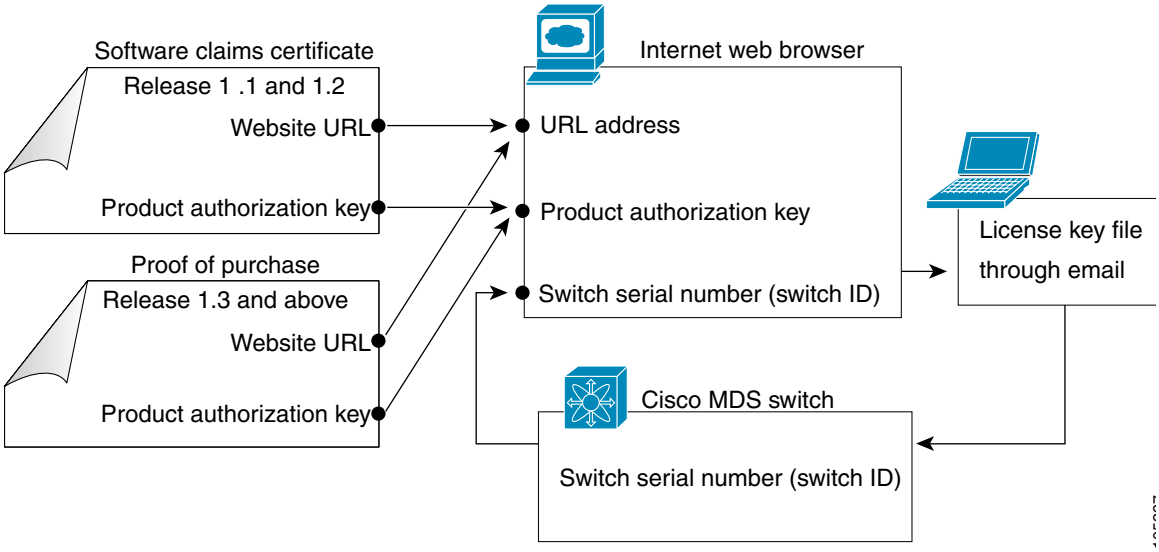

Obtaining License Key Files

To obtain new or updated license key files, follow these steps.

Note The host ID is also referred to as the switch serial number.

Step 1 Obtain your Claim Certificate or the Proof of Purchase document.

This document accompanies every Cisco MDS switch.

Step 2 Locate the Product Authorization Key (PAK) from the Claim Certificate or Proof of Purchase document.

Step 3 Locate the website URL from the Claim Certificate or Proof of Purchase document.

Step 4 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 switch for which it was requested. The requested features are also enabled once the SAN-OS software on the specified switch access the license key file.

Caution Install the license 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 license starts from the first time you start using a feature offered by that license.

Installing Licenses

Note 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.

If you have purchased a new switch through either your reseller or through Cisco, you can have the licenses pre-installed in the factory, or you can install the licenses yourself. If you already have an existing switch, you install the licenses yourself. 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 5-2 for a list of features requiring licenses.

You can install licenses two ways:

- Installing Licenses Using Fabric Manager License Wizard
- Installing Licenses 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 in to 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 in to each switch; however, the switches must be in the fabric you are viewing.

Step 2 Start the License Wizard by clicking on the **License Wizard** icon in the Fabric Manager toolbar (see Figure 5-2); or,

Choose **Switches > License Manager** from the Physical pane. The license information is displayed in the Information pane, one line per feature. Click the **License Keys** tab, and then click the **License Install Wizard** button in the toolbar.

You see the initial screen of the License Wizard.

Figure 5-2 License Install Wizard Icon

Step 3 Select the vendor from which you purchased your switch.

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.

Step 4 If you have already obtained the license key files, click that radio button. Otherwise, click "**I have the Product Authorization Key (PAK)**" if you have the authorization key.

Step 5 Click **Next** to continue to the next screen.

Step 6 Choose the switches for which you have PAKs.

When you check the check box for a switch, the PAK field for that switch becomes editable. The VDH=<serial number> for each switch is shown in the HostId column.

Step 7 Enter the PAK for each switch you have selected.

Step 8 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/uninstall operation completed successfully
- inProgress—License install/uninstall operation is in progress
- corruptedLicenseFile—License file content is Invalid/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 exist with a different expiration date for the feature
- invalidLicenseCount—License count is invalid for the feature

Step 9 Click **Close** to close the wizard. To install more licenses at this point, you must close the wizard and launch it again.

Installing Licenses Using Device Manager

To install a license on your switch using Device Manager, follow these steps:

Step 1 Choose **Admin > License Manager**.

You see the License Manager dialog.

Step 2 Click the **Install** tab to display the Install fields.

The HostId shows the "VDH=" portion of the serial number. The rest of the number is filled in when you complete Steps 3 through 5.

Step 3 Enter the URI from which the license file will be picked for installation.

You should already have copied the license file provided by CISCO-CCO by some other means (for example, through the CLI) to this location.

Step 4 Enter the Target Filename with which the license file will be installed.

Step 5 Click **Install**.

The status of the installation is displayed at the bottom of the dialog box, as follows:

- success—Install/uninstall operation completed successfully
- inProgress—License install/uninstall operation is in progress
- corruptedLicenseFile—License file content is Invalid/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 exist with a different expiration date for the feature
- invalidLicenseCount—License count is invalid for the feature

Step 6 Repeat Steps 3 through 5 to install another license, or click **Close** to close the License Manager dialog.

Viewing License Information in Fabric Manager

To view license information in Fabric Manager, perform these steps:

Step 1 Choose **Switches > License Manager** from the Physical pane. The license information is listed in the Information pane, one line per feature.

Step 2 Click the **Feature Usage** tab to see the Switch, name of the feature package, the type of license installed, the number of licenses used (Usage 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 **File** tab to display the information about each of the License Key Files installed on your switch.

Caution Once an expiration period has started, notifications about license expiration 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.

Viewing License Information in Device Manager

To view license information in Device Manager, perform these steps:

Step 1 Choose **Admin > License Manager**. You see the License Manager dialog.

Step 2 Click the **Features** tab to see the name of the feature package, the type of license installed, the number of licenses used (Usage 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 **Files** tab to display the information about each of the License Key Files installed on your switch.

Removing 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. Features turned on by a permanent license must be disabled before the license in uninstalled.

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.

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 Uninstalling a license requires the related features to first be disabled.

To remove a license, perform the following procedure:

Step 1 Log in to 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 Fabric Manager, select **Switches > License Manager** from the Physical pane. The license information is listed in the Information pane, one line per feature.

From Device Manager, select **Admin > License Manager**. You see the License Manager dialog.

Step 3 In Fabric Manager, click the **File** tab. You see a list of License Key Files. Click on the name of the license you want to remove, and press the **Delete** key or click on 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.

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.

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco, contact Cisco Technical Support at this URL:

License Expiry Alerts

The SAN-OS license counter keeps track of all licenses on a switch. Once an expiry period has started, you will receive CLI console messages, SNMP traps, syslog error messages, and Call Home messages on a daily basis.

Beyond that, the frequency of these message will increase to an hourly basis during the last seven days of the expiry time span. For example:

Your FICON license feature is scheduled to expire in 60 days. If today is December 1st, the license expires on January 30th. In this case, you will receive:

- Daily alerts from December 1st to January 23rd
- Hourly alerts from January 24th to January 29th
- From January 30th, the FICON feature will run without a license for a grace period of 60 days.
- From January 30th to March 21st, you will receive daily alerts about the grace period usage.
- From March 22nd to March 30th, you will receive hourly alerts about the grace period ending.
- On March 31st, the FICON feature is automatically turned off.

License expiry alerts cannot be configured.

Caution After the final seven days of the grace period, the feature is turned off and your network traffic may be disrupted. The grace period also applies to licensed features in Release 1.2(x). While Release 1.2(x) did not enforce the licenses, any upgrade will enforce license requirements and the 60-day grace period.

Moving Licenses 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.

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

Initial Configuration

In order for Cisco MDS 9000 Family switches to be accessed by other devices, they must be initially configured. NTP information is part of this. For the rest of the information, refer to the "Initial Configuration" chapter in the *Cisco MDS 9000 Family Configuration Guide*.

This chapter contains the following topics:

- NTP Configuration, page 6-1
- Display General NTP Statistics for a Switch, page 6-3
- Create an NTP Server or Peer, page 6-3
- Edit an NTP Server or Peer Configuration, page 6-4
- Delete an NTP Server or Peer, page 6-4

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 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 will be 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.

Note 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) act 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.

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).
- Though a peer configured alone, will be the most accurate peer taking on the role of a server, the configured peer should be used more as a back-up support. If more than one server is present, you can have several switches point to one server, and the remaining to the another server, and then configure peer association between these two sets. This forces the clock more reliable.
- If you only have one server, it's better for all the switches have a client association with that server.

If the network is configured robustly, even a server down time will not affect well-configured switches in the network. displays a network with two NTP stratum 2 servers and two switches.

Figure 6-1 NTP Peer and Server Association

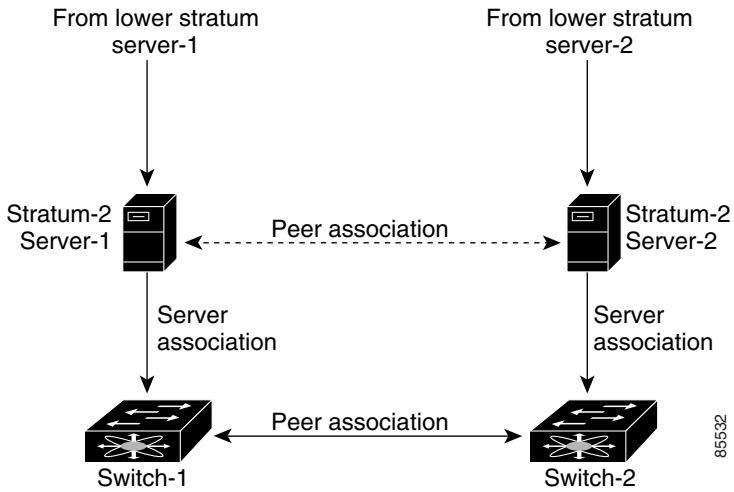

In this configuration, the switches were configured as explained below:

Stratum-2 Server-1	IP address 10.10.10.10	**Switch-1**	IP address 10.10.10.1 NTP configuration: • NTP server 10.10.10.10 • NTP peer 10.10.10.2
Stratum-2 Server-2	IP address 10.10.10.9	**Switch-2**	IP address 10.10.10.2 NTP configuration: • NTP server 10.10.10.9 • NTP peer 10.10.10.1

Display General NTP Statistics for a Switch

To display general NTP statistics for a switch, perform the following steps.

Step 1 From the Fabric Manager Physical pane of the menu tree, choose **Switches > 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 General** tab.

You see the general NTP statistics for that switch.

Create an NTP Server or Peer

To create an NTP server or peer, perform the following steps.

Step 1 From the Fabric Manager Physical pane of the menu tree, choose **Switches > 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 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 **PrefPeer** 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 newly created peer or server is listed on the Peer tab.

Edit an NTP Server or Peer Configuration

To edit an NTP server or peer, perform the following steps.

Step 1 From the Fabric Manager Physical pane of the menu tree, choose **Switches > 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 To change the peer address, double click on the IP address in the Peer Address column, and change the numbers. Alternatively, you can triple click on the IP address and type in a new address.

Step 4 To change the switch mode from peer to server, click on the mode in the Mode column next to the address of the switch.

You see a drop-down list the options peer or server. Select the mode you want for the switch.

Step 5 To change the peer status of the switch to Preferred Peer, check 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, perform the following steps.

Step 1 From the Fabric Manager Physical pane of the menu tree, choose **Switches > 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 To delete a server or peer, click on 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.

Configuring High Availability

This chapter provides details on the high availability feature that is available on switches with two supervisor modules.

This chapter contains the following topics:

- About High Availability, page 7-1
- Switchover Mechanisms, page 7-2
- Process Restartability, page 7-2
- Synchronizing Supervisor Modules, page 7-2
- HA Redundancy States, page 7-2

About 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. The high availability (HA) software framework provides for the following:

- Ensures nondisruptive software upgrade capability.
- 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 Cisco MDS 9216 switches and in the Cisco MDS 9100 Series.
- Protects against link failure using the PortChannel (port aggregation) feature. This feature is also available in Cisco MDS 9216 switches and in the Cisco MDS 9100 Series.
- Provides management redundancy using Virtual Router Redundancy Protocol (VRRP). This feature is also available in Cisco MDS 9216 switches and in the Cisco MDS 9100 Series.
- Switchability--When 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 in the two center slots (sup-1 and sup-2). 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 module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.

Switchover Mechanisms

When the active supervisor module fails, the standby module automatically takes over. You can also 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.

Caution If the supervisor modules are not in a stable state (online or powered down), a switchover will not be performed.

HA Switchover

When a standby supervisor module is in the HA-standby state, an HA switchover is possible. An HA switchover has the following characteristics:

- Is stateful (nondisruptive) since control traffic is not impacted
- Does not impact data traffic since the switching modules are not impacted
- Switching modules are not reset

Process Restartability

Process restartability provides the high availability functionality in Cisco MDS 9000 Family switches.

It ensures that the 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.

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.

HA Redundancy States

The following conditions identify when automatic synchronization is possible:

- If the internal state of one supervisor module is Active with HA standby and of 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 7-1 lists the possible values for the redundancy states.

Table 7-1 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	This module is the active supervisor module and the switch is ready to be configured.
Standby	This state indicate that 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 switch is intentionally shut down for debugging purposes.
At BIOS	The module 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 7-2 lists the possible values for the Supervisor state.

Table 7-2 Supervisor States

State	Description
Active	This module is the active supervisor module and the switch is ready to be configured.
HA standby	This state indicate that 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 7-3 lists the possible values for the internal redundancy state of the supervisor modules.

Table 7-3 Internal States

State	Description
HA standby	This module is the standby supervisor module and the HA switchover mechanism is enabled.
Active with no standby	This state indicate that a switchover is possible.
Active with HA standby	This module is the active supervisor module and 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 supervisor modules.
Standby (failed)	The standby supervisor module is not functioning.
Active with failed standby	This module is 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.

Software Images

This chapter describes how to install and upgrade software images. 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. In the dual supervisor scenario, the standby supervisor module should be updated first.

This chapter contains the following topics:

- About Software Images, page 8-1
- Essential Upgrade Prerequisites, page 8-2
- Using the Software Install Wizard, page 8-3
- Maintaining Supervisor Modules, page 8-4
- Replacing Modules, page 8-5
- Recovering a Corrupted Bootflash, page 8-5
- Default Factory Settings, page 8-5

About Software Images

Each switch is shipped with a Cisco MDS SAN-OS operating system for Cisco MDS 9000 Family switches. The 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 which 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.

Unless explicitly stated, the software install procedures in this section apply to any switch in the Cisco MDS 9000 Family.

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.

  If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco, contact Cisco Technical Support at this URL:

- 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 will be 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 modules or bootflash: (internal to the switch). Ensure that the required free space is available for the image files to be copied.

- 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 IP address for the 10/100 BASE-T Ethernet port connection (interface mgmt0).
  - Ensure that the switch has a route to the remote server. If you do not have a router to route traffic between subnets, the switch and the remote server must be in the same subnetwork. Verify connectivity to the remote server.

- Images
  - The specified system and kickstart images must be 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.
  - Images can be retrieved in one of two ways:

    Local—Images are locally available on the switch.

    Remote—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.

Table 8-1 summarizes terms used in this chapter with specific reference to the install and upgrade process.

Table 8-1 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 non-disruptive).

Table 8-1 Terms Specific to this Chapter (continued)

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.

Using the Software Install Wizard

Note Before you use this wizard, be sure the standby supervisor management port is connected.

To use the Software Install Wizard, perform the following steps.

Step 1 Open the Software Install Wizard by clicking on its icon in the toolbar (see Figure 8-1).

Figure 8-1 Software Install Wizard Icon

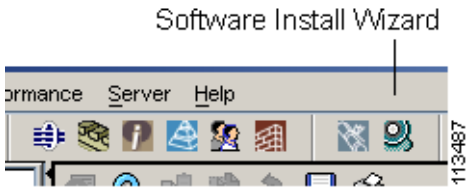

You see the Software Install Wizard.

Step 2 To manage switch images, choose the switches from the list.

You must select at least one switch in order to proceed. When finished, click **Next**.

Step 3 For each switch model, specify the new images to use.

Click **...** to enter image URIs and other information. You must specify at least one image for each switch in order to proceed. The total space required on the bootflash to copy the image is shown in the Required Flash Space column.

To use images that are already downloaded (the file is already on the bootflash), check the **Skip Image Download** check box. When you check this check box, you are prompted to choose an image from the bootflash for each switch being upgraded.

Step 4 Check the active (and standby, if applicable) bootflash on each switch to see if there is enough space for the new images.

The table 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. Free additional space by clicking **...**, or deselect the switch by going back to the first screen and unchecking the check box.

Step 5 For each switch, click **...** to choose images from the bootflash to use for the upgrade.

You must choose at least one image for each switch in order 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 6 Start the upgrade.

If you check "version check" before the upgrade process is started, a version check is done. This 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, prompting you to confirm that this check should be performed.

Caution If "version check" is enabled, the upgrade will proceed even if your version is newer than the version you are installing.

Note On hosts where the TFTP server can not be started, a warning displays. 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 displays a status during processing. Check the lower left-hand corner of the wizard for the status message "Upgrade Finished." First, the wizard displays the message "Success 1.2(2)" followed a few seconds later by "InProgress Polling." Then the wizard displays a second message "Success 1.2(2)" before displaying the final "Upgrade Finished."

Maintaining Supervisor Modules

This section includes general information about replacing and using supervisor modules effectively.

Standby Supervisor Boot Variable Version

If the standby supervisor module's 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 Boot Alert

If a standby supervisor module fails to boot, the active supervisor module detects that conditions and generates a Call Home event and a SYSLOG message and reboots the standby supervisor module approximately 3 to 6 minutes after the standby supervisor module moves to the loader> prompt.

The following SYSLOG error 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 the *Cisco MDS 9000 Family San Volume Controller Configuration Guide* for configuration details on replacing the Caching Services Module (CSM).

When a spare standby supervisor module is inserted, it uses the same image as the active supervisor module. The SAN-OS software image is not copied to the standby flash until you issue an **install all** command. Refer to the *Cisco MDS 9000 Family Configuration Guide* for more information.

Recovering a Corrupted Bootflash

All switch configurations reside in the internal bootflash. If you have a corrupted internal bootflash, you could potentially lose your configuration. Be sure to save and back up your configuration files periodically. Refer to the *Cisco MDS 9000 Family Configuration Guide* for information on recovering a corrupted bootflash.

Default Factory Settings

Table 8-2 lists the default settings for all Cisco MDS 9000 Family switches.

Table 8-2 Default Factory Settings

Parameters	Default
Kickstart image	No image is specified
System image	No image is specified

Managing Modules

This chapter describes how to manage switching modules (also known as line cards) and provides information on monitoring module states.

This chapter contains the following topics:

- About Modules, page 9-1
- Viewing the State of a Module, page 9-2
- Identifying Module LEDs, page 9-3
- Configuring EPLDs, page 9-5
- Default Supervisor Module Settings, page 9-6

About Modules

Table 9-1 describes the supervisor module options for switches in the Cisco MDS 9000 Family.

Table 9-1 Supervisor Module Options

Product	No. of Supervisor Modules	Supervisor Module Slot	Switching Module Features
Cisco MDS 9216	One module (includes 16 Fibre Channel ports)	1	2-slot chassis allows one optional switching module in the other slot.
Cisco MDS 9509	Two modules	5 and 6	9-slot chassis allows any switching module in the other seven slots.
Cisco MDS 9506	Two modules	5 and 6	6-slot chassis allows any switching module in the other four slots.

Supervisor Modules

Supervisor modules are automatically powered up and started with the switch.

Cisco MDS 9200 Series switches have one supervisor module that includes an integrated 16-port switching module.

Cisco MDS 9500 Series switches have two supervisor modules--one in slot 5 (sup-1) and one in slot 6 (sup-2). When the switch powers up and both supervisor modules come up together, the module that enters the active mode is dependent on which of the two modules 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.

Switching Modules

Cisco MDS 9000 Family switches support any switching module in any non-supervisor slot. The switching module obtains its image from the supervisor module.

The interfaces in each module are ready to be configured when the module displays an ok status.

Viewing the State of a Module

The switching module goes through a testing and an initializing stage before displaying an ok status. Table 9-2 describes the possible states in which a module can exist.

Table 9-2 Module States

show module Output	Description
`powered up`	The hardware has electrical power. When the hardware is powered up, the software begins booting.
`testing`	The module has established connection with the supervisor 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`	This module is the standby supervisor module and that the HA switchover mechanism is enabled.
`standby`	This module is the standby supervisor module and the warm switchover mechanism is enabled.

To view the state of a module from Device Manager, choose **Physical > Modules**. The dialog box displays the status of every module.

Identifying Module LEDs

Table 9-3 describe the LED location, type, and status for supervisor and switching modules used in Cisco MDS 9000 Family switches.

Table 9-3 Module LEDs on a Cisco MDS 9200 Series Switch

Module	LED Type	Status	Description
Fixed switching module	Status	Green	• All chassis environmental monitors (power supply, fan, temperature sensor, clock, and chassis) are reporting OK. • Sufficient power is available for all modules
		Orange	• Any one of the chassis environmental monitors (power supply, fan, temperature sensor, clock, and chassis) failed. • Sufficient power is not available for all modules. • Incompatible power supplies are installed. • The redundant clock failed.
		Red	• The diagnostic test failed. • The module is not operational because a fault occurred during the initialization sequence. • A temperature condition occurred. (A major threshold was exceeded during environmental monitoring.)
Optional switching module	System	Green	All diagnostics pass. The module is operational (normal initialization sequence).
		Orange	• The module is booting or running diagnostics (normal initialization sequence). • An over temperature condition occurred. (A minor threshold was exceeded during environmental monitoring.)
		Red	• 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.)

Table 9-4 lists the system and power management LEDs on a redundant supervisor module that are synchronized to the active supervisor module.

Table 9-4 Supervisor Module LEDs on a Cisco MDS 9500 Series Switch

LED	Status	Description
Status	Green	All diagnostics pass. The module is online.
	Orange	• The module is booting or running diagnostics (normal initialization sequence). • The module is not online. • An over temperature condition has occurred. (A minor threshold has been exceeded during environmental monitoring.)
	Red	• The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. • An over temperature condition has occurred. (A major threshold has been exceeded during environmental monitoring.)
System	Green	All chassis environmental monitors (power supply, fan, temperature sensor, clock, and chassis) are reporting OK.
	Orange	• Any one of the environmental monitors (power supply, fan, temperature sensor, clock, and chassis) 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 Mgmt1	Green	Sufficient power is available for all modules.
	Orange	Sufficient power is not available for all modules.

Table 9-5 lists the Ethernet interface LEDs on a Cisco MDS 9200 Series Switch.

Table 9-5 Ethernet Interface LEDs on a Cisco MDS 9200 Series Switch

Module	LED Type	Status	Description
Ethernet (mgmt 0)	Activity	Flashing green	Traffic is passing through the interface.
	Link	Solid green	The link is functioning.
		Off	The link is down.

Table 9-6 lists the switching module LEDs.

Table 9-6 Switching Module LEDs

LED Type	Status	Description
Status	Green	All diagnostics pass. The module is operational (normal initialization sequence).
	Orange	• The module is booting or running diagnostics (normal initialization sequence). • An over temperature condition occurred. (A minor threshold was exceeded during environmental monitoring.)
	Red	• 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.)
Speed	On	2 Gbps mode.
	Off	1 Gbps mode.
Link	Solid green	Link is up.
	Flashing green	Link is up (beacon used to identify port).
	Solid yellow	Disabled by software.
	Flashing yellow	Fault is detected.
	Off	Link is down.

Configuring EPLDs

Switches and directors in the Cisco MDS 9000 Family contain several electrically programmable logical devices (EPLDs) that provide hardware functionalities in all modules. Starting with Cisco MDS SAN-OS Release 1.2, EPLD image upgrades are periodically provided to include enhanced hardware functionality or to resolve known issues.

Tip Refer to the *Cisco MDS SAN-OS Release Notes* to verify if the EPLD has changed for the SAN-OS image version being used.

EPLDs can be upgraded or downgraded. When EPLDs are being upgraded or downgraded, the following guidelines and observations apply:

- You can individually update each module that is online. The EPLD update is only disruptive to the module being upgraded.
- If you interrupt an upgrade, the module must be upgraded again.
- The upgrade or downgrade can only be executed from the active supervisor module. While the active supervisor module cannot be updated, you can update the other modules individually.
- In Cisco MDS 9100 Series Fabric switches, be sure to specify 1 as the module number.
- Cisco MDS 9216 Switches do not support EPLD upgrades.

Caution Do not insert or remove any modules while an EPLD upgrade or downgrade is in progress.

Note Switches in the Cisco MDS 9100 Series do not support a forced EPLD upgrade. When you upgrade the EPLD module on these switches, you receive the following message:

```
Data traffic on the switch will stop now!! Do you want to continue (y/n)?
```

Refer to the *Cisco MDS 9000 Family Configuration Guide* for information on upgrading EPLDs.

Default Supervisor Module Settings

Table 9-7 lists the default settings for the supervisor module.

Table 9-7 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 is set to 0.0.0.0.

Managing System Hardware

This chapter provides details on monitoring the health of the switch.

This chapter contains the following topics:

- Configuring Power Supplies, page 10-1
- Displaying Module Temperature, page 10-5
- Monitoring Fan Modules, page 10-6
- Monitoring Clock Modules, page 10-6
- Viewing System Attributes, page 10-6
- Viewing Running Processes, page 10-6
- Viewing Flash File Information, page 10-7
- Managing Inventory Information, page 10-7
- Managing Module Attributes, page 10-7

Configuring Power Supplies

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 Watts when powered at 110 volts, and 2500 Watts when powered at 220 volts.

Guidelines for Power Supplies with Different Capacities

When power supplies with different capacities are installed in the switch, the total power available differs based on the configured mode.

- In redundant mode, the total power is the lesser of the two power supply capacities. For example, if you have the following usage figures configured:

  Power supply 1 = 2500 Watts

  Additional Power supply 2 = not used

  Current usage = 2000 Watts

  Current capacity = 2500 Watts

  Then the following three scenarios differ as specified:

  – **Scenario 1**—If 1800 Watts is added as power supply 2, then power supply 2 is shut down.

    Reason: 1800 Watts is less than the usage of 2000 Watts.

  – **Scenario 2**—If 2200 Watts is added as power supply 2, then the current capacity decreases to 2200 Watts.

    Reason: 2200 Watts is the lesser of the two power supplies.

  – **Scenario 3**—If 3000 Watts is added as power supply 2, then the current capacity value remains at 2500 Watts.

    Reason: 2500 Watts is the lesser of the two power supplies.

Table 10-1 lists redundant mode power supply scenarios.

Table 10-1 Redundant Mode Power Supply Scenarios

Scenario	Power Supply 1 (Watts)	Current Usage (Watts)	Insertion of Power Supply 2 (Watts)	New Capacity (Watts)	Action Taken by Switch
1	2500	2000	1800	2500	Power supply 2 is shut down.
2	2500	2000	2200	2200	Capacity becomes 2200Watts.
3	2500	2000	3300	2500	Capacity remains the same.

- In combined mode, the total power is twice the lesser of the two power supply capacities.

  For example, if you have the following usage figures configured:

  Power supply 1 = 2500 Watts

  Additional Power supply 2 = not used

  Current Usage = 2000 Watts

  Current capacity = 2500 Watts

  Then, the following three scenarios differ as specified:

  – **Scenario 1**—If 1800 Watts is added as power supply 2, then the capacity increases to 3600 Watts.

    Reason: 3600 Watts is twice the minimum (1800 Watts).

  – **Scenario 2**—If 2200 Watts is added as power supply 2, then the current capacity increases to 4400 Watts.

Reason: 4400 Watts is twice the minimum (2200 Watts).

– **Scenario 3**—If 3000 Watts is added as power supply 2, then the current capacity increases to 5000 Watts.

Reason: 5000 Watts is twice the minimum (2500 Watts).

Table 10-2 lists combined mode power supply scenarios.

Table 10-2 Combined Mode Power Supply Scenarios

Scenario	Power Supply 1 (Watts)	Current Usage (Watts)	Insertion of Power Supply 2 (Watts)	New Capacity (Watts)	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	

Guidelines for Power Supplies with Different Capacities

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. Various configuration scenarios are displayed and summarized in Table 10-3.

- **Scenario 1**—You have the following usage figures configured:

  Power supply 1 = 2500 Watts

  Additional Power supply 2 = 1800 Watts

  Current Usage = 2000 Watts

  Current mode = **combined** mode (so current capacity is 3600Watts)

  You decide to change the switch to **redundant** mode. Then power supply 2 is shut down.

  Reason: 1800Watts 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 Watts

  Additional Power supply 2 = 2200 Watts

  Current Usage = 2000 Watts

  Current mode = **combined** mode (so current capacity is 4400Watts).

  You decide to change the switch to **redundant** mode. Then the current capacity decreases to 2200Watts.

  Reason: 2200Watts is the lesser of the two power supplies.

- **Scenario 3**—You have the following usage figures configured:

  Power supply 1 = 2500 Watts

  Additional Power supply 2 = 1800 Watts

  Current Usage = 3000 Watts

  Current mode = **combined** mode (so current capacity is 3600 Watts).

You decide to change the switch to **redundant** mode. Then the current capacity decreases to 2500Watts and the configuration is rejected.

Reason: 2500Watts is less than the system usage (3000Watts).

Table 10-3 Summary of Configuration Scenarios for Power Supplies with Different Capacities

Scenario	Power Supply 1 (W)	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	Existing configuration.
	2500	**N/A**	2000	1800	**redundant**	2500	Power supply 2 is shut down
2	2500	**combined**	2000	2200	**N/A**	4400	Existing configuration.
	2500	**N/A**	2000	2200	**redundant**	2200	The new capacity is changed.
3	2500	**combined**	3000	1800	**N/A**	3600	Existing configuration.
	2500	**N/A**	3000	1800	**redundant**	N/A	Rejected, so the mode reverts to **combined** mode.

Managing Power Supplies

To manage power supply power attributes, perform the following steps.

Step 1 From the Fabric Manager menu tree, choose **Switches > Modules** and then click the **Power Supplies** tab, or from the Device Manager, choose **Physical > Power Supplies**. The Information pane from the Fabric Manager displays power supply power attributes for multiple switches. The dialog box from the Device Manager displays power supply power attributes for a single switch.

Step 2 Configure the power attributes for the power supply.

Displaying Module Temperature

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 Celsius), minor and major.

Note A threshold value of -127 indicates that no thresholds are configured or applicable.

When a minor threshold is exceeded, a minor alarm occurs and the following action is taken for all four sensors:

- Syslog messages are displayed.
- Call Home alerts are sent (if configured).
- SNMP notifications are sent (if configured).

When a major threshold is exceeded, a major alarm occurs and the following action is taken as follows:

- For sensors 1, 3, and 4 (outlet and onboard sensors):
  - Syslog 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, the module is shut down.
  - If the threshold is exceeded in a supervisor module with HA-standby or standby present, the supervisor module is shut down.
  - If the standby supervisor is not present, the entire switch is shut down.

Note Switch shut down only happens after a two-minute interval. During this interval the software monitors the temperature every five (5) seconds and continuously sends syslog messages as configured. If the required action is not taken (for example, a new fan module is inserted to decrease temperature) and if the temperature does not come down, the system is shut down at the end of two minutes.

To monitor sensor temperature attributes, perform the following steps.

Step 1 From the Fabric Manager menu tree, choose **Switches > Modules** and then click the **Temperature Sensors** tab, or from Device Manager, choose **Physical > Temperature Sensors**.

The Information pane from the Fabric Manager displays sensor temperature attributes for multiple switches. The Sensors dialog box from the Device Manager displays sensor temperature attributes for a single switch.

Step 2 Configure the sensor attributes.

Monitoring Fan Modules

The fan status is continuously monitored. In case of a fan module failure, the following action is taken:

- Syslog messages are displayed.
- Call Home alerts are sent (if configured).
- SNMP notifications are sent (if configured).

Caution A fan failure could lead to temperature alarms if not corrected immediately.

When a fan module is removed, the Fan module removed message is displayed every 10 seconds for three minutes after the fan removal. After the three-minute interval, a system shutdown message is displayed every 5 seconds for two additional minutes. At the end of this 5-minute sequence, the system is shutdown. If the fan module is re-inserted at any point within this 5-minute period, the remaining sequence is stopped.

To monitor a fan module from Device Manager, choose **Physical > Fans**. The dialog box displays the fan status.

Monitoring Clock Modules

Each switch has two clock modules for redundancy: Clock A (primary) and Clock B. The redundant clock module (Clock B) takes over if the primary clock module fails. If Clock A is available at startup, the switch uses Clock A, otherwise it uses Clock B.

If Clock A fails, the switch is reset and Clock B automatically takes over. Clock modules cannot be configured. If both modules fail, the switch shuts down. The probability of a clock failure is low given that the mean time between failures (MTBF) is 3660316 hours.

Viewing System Attributes

To manage system attributes, perform the following steps.

Step 1 From the Fabric Manager menu tree, choose **Switches**, or from Device Manager, choose **Admin > System**. You can also double-click on the chassis in Device Manager.

The Fabric Manager Information pane shows system attributes for multiple switches. The dialog box from the Device Manager shows system attributes for a single switch.

Step 2 Configure the system attributes for the chassis.

Viewing Running Processes

You can view information about the processes that are currently running on a switch from Device Manager. Choose **Admin > Running processes** to see the process information listed in the Running Processes dialog box.

Viewing Flash File Information

You can view information about the files that are currently stored in flash memory on the switch from Device Manager. Choose **Admin > Flash Files** to see the stored files information listed in the Flash Files dialog box.

Managing Inventory Information

To manage inventory attributes, perform the following steps.

Step 1 From the Fabric Manager menu tree, choose **Switches > Modules**, and then click the **Inventory** tab, or from Device Manager, choose **Physical > Inventory**.

The Fabric Manager Information pane displays system attributes for multiple switches. The dialog box from the Device Manager displays system attributes for a single switch.

Step 2 Configure the inventory attributes for the module.

Managing Module Attributes

To manage module status attributes, perform the following steps.

Step 1 From the Fabric Manager menu tree, choose **Switches > Modules** and then click the **Card Status** tab, or from Device Manager, choose **Physical > Modules**.

The Information pane from the Fabric Manager displays card attributes for multiple switches. The dialog box from the Device Manager view displays attributes for a single switch.

Step 2 Configure the status attributes for the module.

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 which allows identical Fibre Channel IDs (FCIDs) to be used simultaneously in different VSANs. 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 is provided (to another VSAN in the same physical SAN) by a configured 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.

This chapter contains the following topics:

- How VSANs Work, page 11-2
- VSANs Versus Zones, page 11-4
- Default and Isolated VSANs, page 11-5
- VSAN Membership, page 11-6
- VSAN Attributes, page 11-6
- Adding and Configuring VSANs, page 11-7
- Deleting VSANs, page 11-7
- Default Settings, page 11-8

How VSANs Work

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.

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 (FCIDs) 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.

Figure 11-1 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. Within each VSAN, all members can talk to one another. Between VSANs no communication is possible.

Figure 11-1 Logical VSAN Segmentation

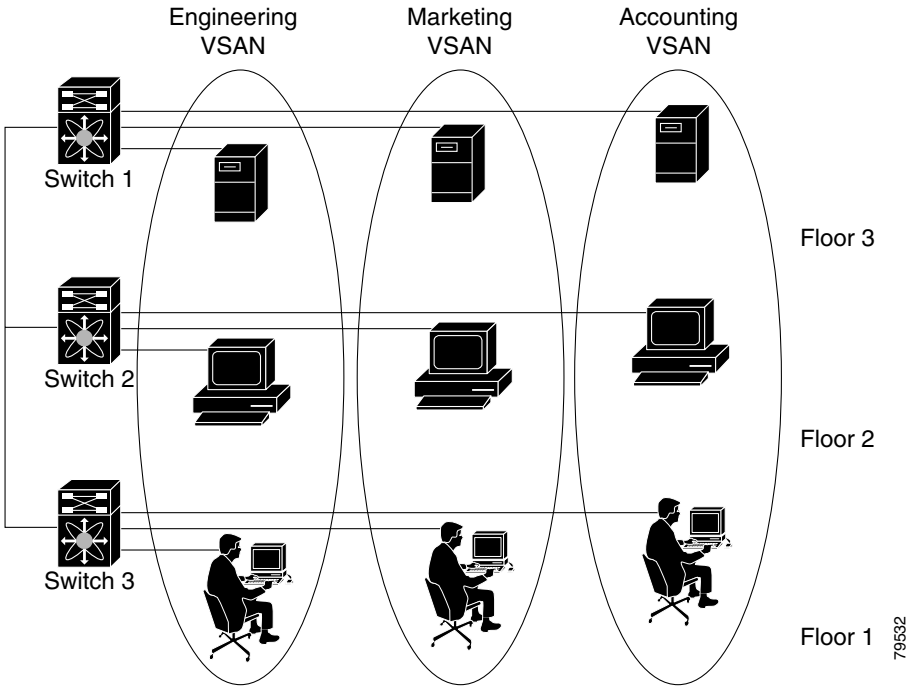

Figure 11-2 shows a physical Fibre Channel switching infrastructure with two defined VSANs: VSAN 2 (dashed) and VSAN 7 (solid). VSAN 2 includes hosts H1 and H2, application servers AS2 and AS3, and storage arrays SA1 and SA4. VSAN 7 connects H3, AS1, SA2, and SA3.

The switch icons shown in both figures indicate that these features apply to any switch in the Cisco MDS 9000 family.

Figure 11-2 Example of Two VSANs

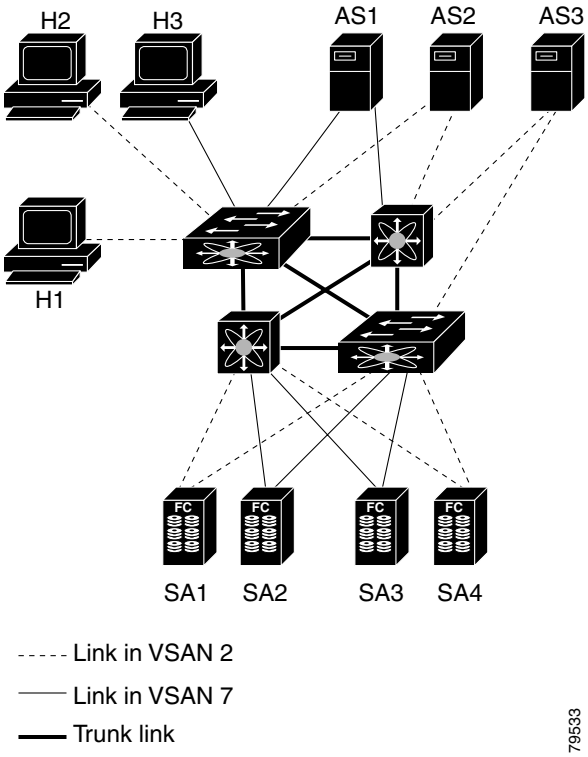

The four switches in this network are interconnected by trunk links that carry both VSAN 2 and VSAN 7 traffic. Thus the inter-switch topology of both VSAN 2 and VSAN 7 are identical. This is not a requirement and a network administrator can enable certain VSANs on certain links to create different VSAN 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. 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.

VSANs Versus Zones

You can define multiple zones in a VSAN. Because two VSANs are equivalent to two nonconnected SANs, zone A on VSAN 1 is different and separate from zone A in VSAN 2. Table 11-1 lists the differences between VSANs and zones.

Table 11-1 VSAN and Zone Comparison

VSANs	Zones
VSANs equal SANs with routing, naming, and zoning protocols.	These 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 may 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 11-3 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 11-3 VSANS with Zoning

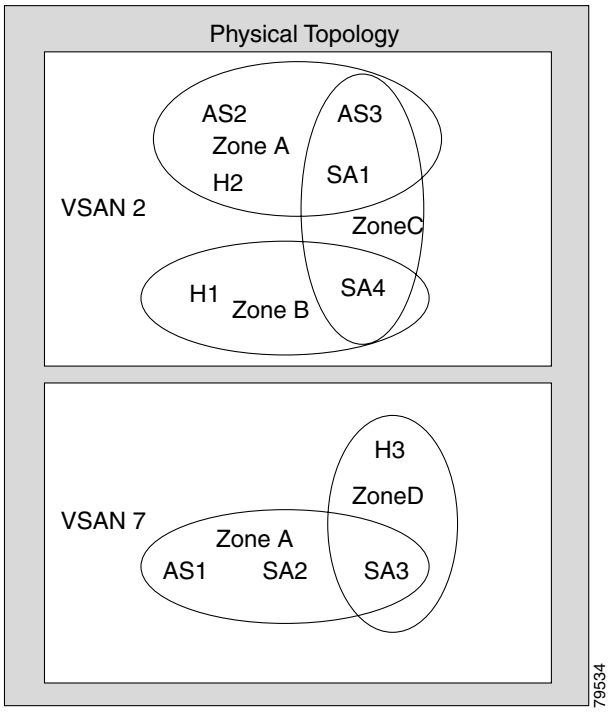

Default and Isolated VSANs

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.

Default VSANs

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. It can be suspended.

Isolated VSANs

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.

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.

Trunking ports have an associated list of VSANs that are part of an allowed list.

VSAN Attributes

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.

- Load balancing attributes—These 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.

Operational State of a VSAN

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.

Adding and Configuring VSANs

To add and configure VSANs, perform the following steps.

Step 1 From the Fabric Manager menu tree in the Logical pane, click on the desired VSAN, or from Device Manager, choose **FC > VSAN**, or click the **VSAN** icon on the toolbar.

The Fabric Manager's Information pane displays VSAN attributes for multiple switches. The VSAN dialog box in the Device Manager displays VSAN general attributes for a single switch.

Step 2 From Fabric Manager, click **Create VSAN**, or from Device Manager, click **Create**.

You see the Create dialog box.

Step 3 Complete the fields on this dialog box and click **Create** to add the VSAN.

Deleting VSANs

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.

- VSAN attributes and port membership details are maintained by 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.

Figure 11-4 VSAN Port Membership Details

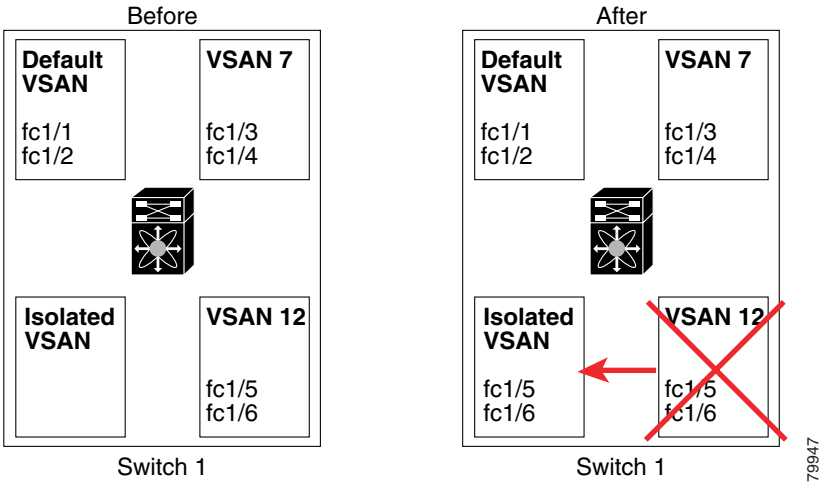

- VSAN-based runtime (name server), zoning, and configuration (static routes) information is removed when the VSAN is deleted.
- Configured VSAN interface information is removed when the VSAN is deleted.

Note The allowed VSAN list is not affected when a VSAN is deleted.

Default Settings

The table below lists the default settings for all configured VSANs.

Table 11-2 Default VSAN Parameters

Parameters	Default
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).
Port membership	Default VSAN (VSAN 1).

Configuring Interfaces

A switch's 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, management interface (mgmt0), or VSAN interfaces.

This chapter describes the basic interface configuration to get your switch up and running.

Before you begin configuring the switch, ensure that the modules in the chassis are functioning as designed.

This chapter contains the following topics:

- Configuring Fibre Channel Interfaces, page 12-1
- Default Settings, page 12-10
- Configuring VSAN Interfaces, page 12-10
- Configuring Gigabit Ethernet Interfaces, page 12-11
- Enabling or Disabling Ports, page 12-11
- Managing Interface Attributes for Ports, page 12-11

Configuring Fibre Channel Interfaces

This section describes Fibre Channel interface characteristics, including (but are not limited to) modes, states, and speeds.

This section contains the following topics:

- About Interface Modes, page 12-2
- About Interface States, page 12-4
- Configuring TL Port ALPA Caches, page 12-8
- Configuring Buffer-to-Buffer Credits, page 12-8
- Configuring Performance Buffers, page 12-8
- Configuring the Beacon Mode, page 12-9
- Identifying the Beacon LEDs, page 12-9
- Configuring Switch Port Defaults, page 12-10

About Interface Modes

Each physical Fibre Channel interface in a switch may operate in one of several modes: E port, F port, FL port, TL port, TE port, and SD port. Besides these modes, each interface may be configured in auto or Fx port mode. These two modes determine the port type during interface initialization. Figure 12-1 shows different interface modes.

Figure 12-1 Cisco MDS 9000 Family Switch Interface Modes

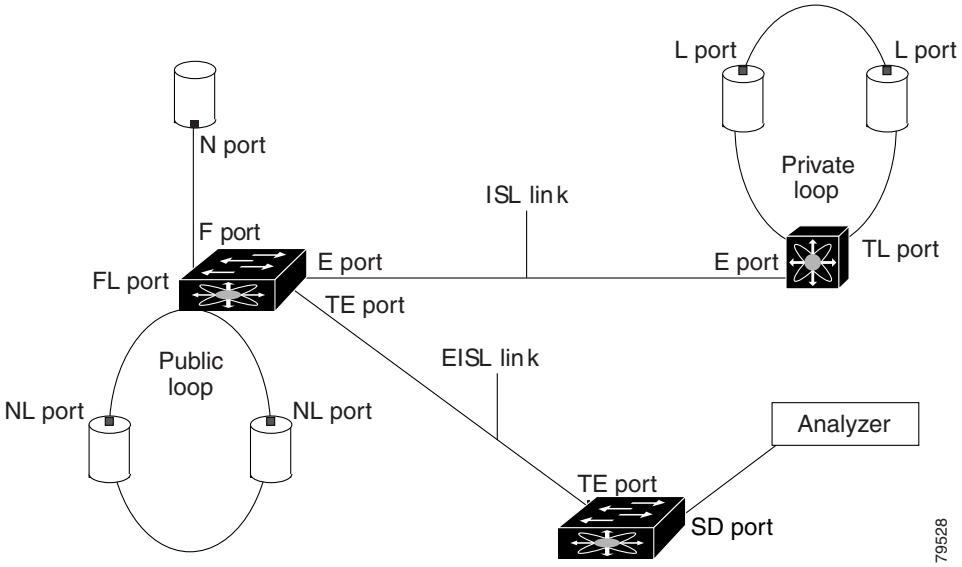

Interfaces are created in VSAN 1 by default.

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).

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

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.

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 port mode is 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.

Note We recommend configuring devices attached to TL ports in zones which have up to 64 zone members.

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 Extended 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. 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.

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.

Fx Port

Interfaces configured as Fx ports are allowed to 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 Cisco's 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.

When an FCIP peer is a SAN extender device that only support 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.

Auto Mode

Interfaces configured as auto are allowed to 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. TL ports and SD ports are not determined during initialization and are administratively configured.

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 the Table 12-1.

Table 12-1 Administrative States

Administrative State	Description
Up	Enables an interface.
Down	Disables an interface. When an interface is administratively disabled, the physical link layer state change is ignored.

Operational States

The operational state indicates the current operational state of the interface. Table 12-2 describes the operational states.

Table 12-2 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. Table 12-3 describes the reason codes for interface states.

Table 12-3 Reason Codes for Interface States

Administrative Configuration	Operational Status	Reason Code
Up	Up	None.
Down	Down	Administratively down--If you administratively configure an interface as down, you disable the interface. No traffic is received or transmitted.
Up	Down	See the following table.

If the administrative state is up and the operational state is down, the reason code differs based on the nonoperational reason code. Table 12-4 describes the reason codes for nonoperational states.

Table 12-4 Reason Codes for Nonoperational States

Reason Code	Description	Applicable Modes
Link failure or not connected	Physical layer link is not operational.	All
Fcot not present	The Fibre Channel optical transmitter hardware (FCOT) 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	Waiting 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 configure the interface as **shutdown** followed by **no shutdown**.	
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.	

Table 12-4 Reason Codes for Nonoperational States (continued)

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.	

Configuring 32-port Switching Modules and Host-Optimized Ports

The 32-port 1/2-Gbps switching module contains 8 port groups of 4 ports each. When configuring these modules or the host-optimized ports in the Cisco 9100 Series, the following 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 in the shutdown state.
- If any of the other three ports are configured in a no shutdown state, you cannot configure the first port as an E port. The other three ports continue to remain in a no shutdown state.
- The default port mode is auto. The Auto option 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).

- 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 TL Port ALPA Caches

While 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 has an ALPA, the 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 SAN-OS software discards an inactive cache entry (if available) to make space for the new entry.

Configuring 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, since switches must not drop frames. Buffer Credits are negotiated on a per-hop basis.

The receive BB_credit (rxbbcredit) value may be configured for each FC interface. In most cases, you don't need to modify the default configuration.

The receive BB_credit values depend on the module type and the port mode:

- 16-port switching modules and full rate ports—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.
- 32-port switching modules and host-optimized ports—The default value is 12 for the 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 Performance Buffers

Regardless of the configured Rx 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 use the Default option, the built-in algorithm is used.

If you do not specify this command, the Default option is automatically used.

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.

Identifying the Beacon LEDs

Figure 12-2 displays the status, link, and speed LEDs in a 16-port switching module.

Figure 12-2 Cisco MDS 9000 Family Switch Interface Modes

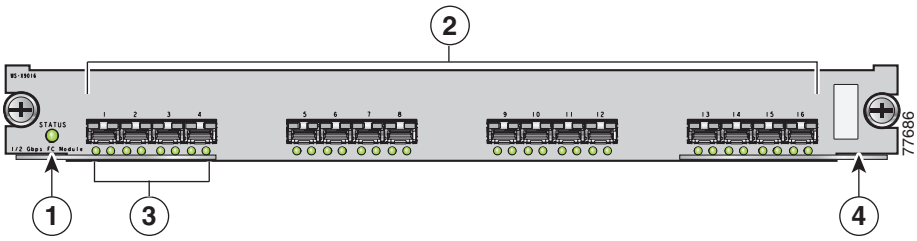

1	Status LED	**3**	Link LEDs and speed LEDs
2	1/2-Gbps Fibre Channel port group	**4**	Asset tag (refer to the *Cisco MDS 9000 Family Hardware Installation Guide*).

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).
- On (flashing green)—The interface is in loopback mode.

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.

Configuring Switch Port Defaults

You can configure default values for various switch port attributes. If you configure the following attributes, they will be applied globally to all future switch port configurations, even if you do not individually specify them at that time.

Default Settings

Table 12-5 lists the default settings for Fibre Channel interface parameters.

Table 12-5 Default Fibre Channel 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

Configuring 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.

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 manually.
- If you delete the VSAN, the attached interface is automatically deleted.

You can configure each interface only in one VSAN.

After configuring the VSAN interface, you can configure an IP address or Virtual Router Redundancy Protocol (VRRP) features.

Configuring Gigabit Ethernet Interfaces

Each port or interface on the IPS module is displayed in the Ethernet Port dialog. To configure Ethernet port interfaces, perform the following steps:

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Open Device Manager.

Step 3 Choose any Ethernet port by clicking on it once.

Step 4 Choose **Interfaces > Gigabit Ethernet Ports**. All Gigabit Ethernet ports for the switch are displayed in a table.

Step 5 To configure the alias, state, or IP address for a particular port, double-click on the appropriate table cell.

Step 6 Enter the alias, IP address, or state for the port.

Step 7 Click **Apply**.

Enabling or Disabling Ports

To enable a port, right-click on a disabled port in Device Manager and choose Enable from the pop-up menu. To disable a port, right-click on a enabled port in Device Manager and choose Disable from the pop-up menu.

To enable or disable multiple ports, press the **Ctrl** key and click each port or drag the mouse around a group of ports. Then right-click any of the chosen ports and choose either Enable or Disable from the pop-up menu.

Managing Interface Attributes for Ports

To manage port interface attributes from the Fabric Manager, choose Physical Interfaces from the menu tree and then choose one of the following port types to be configured:

- Port Channels
- xEPorts
- FxPorts
- Other Ports

To manage port interface attributes from the Device Manager, select a port on a module, and then choose a port type from the Interface menu.

The Fabric Manager Information pane displays interface attributes for multiple switches. The dialog box from Device Manager displays interface attributes for a single switch.

Configuring Trunking

This chapter describes the trunking feature provided in Cisco MDS 9000 switches. It includes the following sections:

This chapter contains the following topics:

- About Trunking, page 13-1
- About Trunking Protocol, page 13-2
- Configuring Trunk Modes, page 13-2
- Configuring Trunk-Allowed VSAN List, page 13-2
- Trunking Configuration Guidelines, page 13-4
- Default Settings, page 13-5

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 Extended ISL (EISL) frame format.

Figure 13-1 Trunking

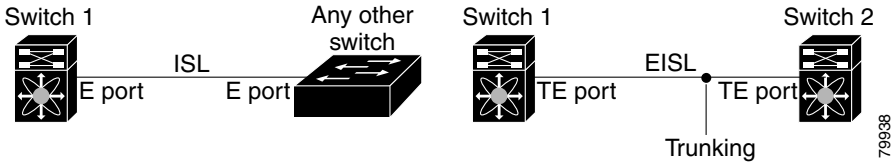

The trunking feature includes the following restrictions:

- 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.

About 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 will not be affected--the TE port continues to function in trunk mode, but only supports traffic in VSANs that it negotiated 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.

To avoid inconsistent configurations, shut all E ports before enabling or disabling the trunking protocol.

Configuring Trunk Modes

By default, the trunk mode is enabled in all Fibre Channel interfaces. However, the trunk mode configuration takes effect only in E-port mode. You can configure the 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 resulting trunking state of the link and the port modes at both ends.

Table 13-1 Trunk Mode Status Between Switches

Your Trunk Mode Configuration	Resulting State and Port Mode
Switch 1	**Switch 2**
On	Auto or on
Off	Auto, on, or off
Auto	Auto

When connected to a third-party switch, the trunk mode configuration has no effect--the ISL is always in a trunking disabled state.

Configuring Trunk-Allowed VSAN List

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 13-2, 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 13-2 Default Allowed -Active VSAN Configuration

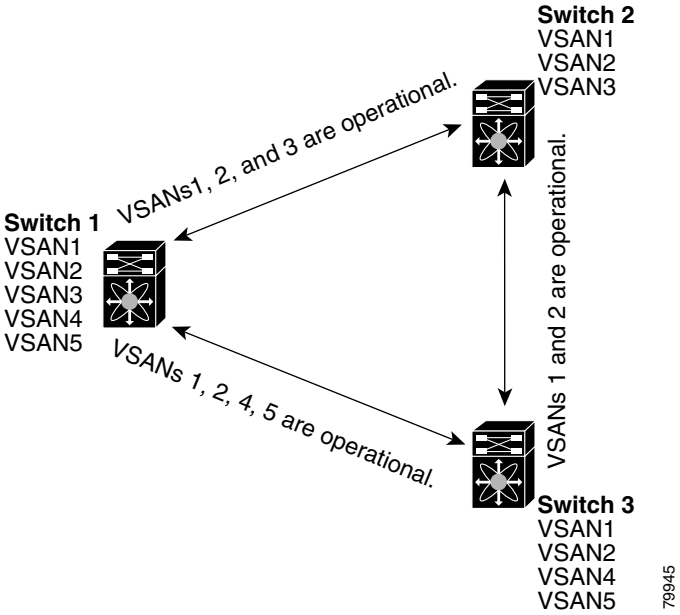

You can configure a select set of VSANs (from the allowed-active list) to control access to those VSANs in a trunking ISL. Using Figure 13-3 as an example, you can configure the list of allowed VSANs on a per-interface basis.

Figure 13-3 Operational and Allowed VSAN Configuration

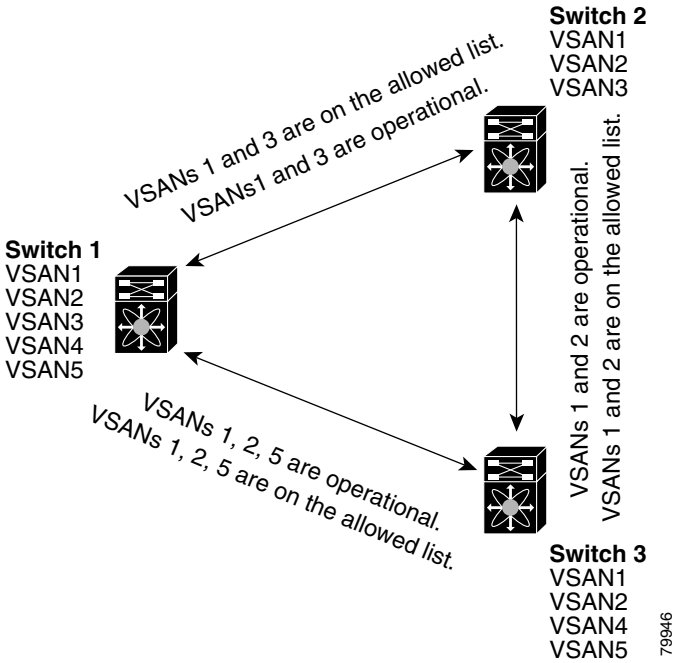

In Figure 13-3, the operational allowed list of VSANs between switches is as follows:

- Switch 1 and switch 2 include VSAN 1 and VSAN 3.
- Switch 2 and switch 3 include VSAN 1 and VSAN 2.
- Switch 3 and switch 1 include VSAN 1, 2, and 5.

Consequently, VSAN 2 can only be routed from switch 1 through switch 3 to switch 2.

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 13-4.)

Figure 13-4 VSAN Mismatch

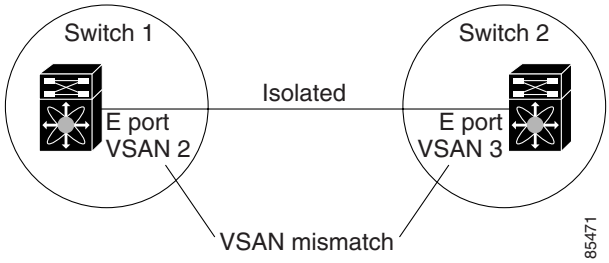

In Figure 13-4, the trunking protocol detects potential VSAN merging and isolates the ports involved.

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 13-5.)

Figure 13-5 Third-Party Switch VSAN Mismatch

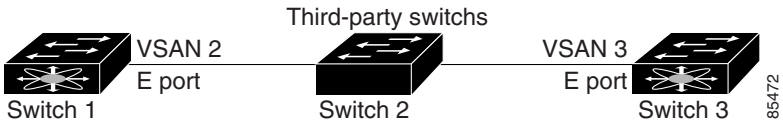

VSANs 2 and 3 get effectively merged with overlapping entries in the name server and the zone applications. The Cisco MDS 9000 Fabric Manager helps detect such topologies. Refer to the *Cisco MDS 9000 Family Fabric Manager User Guide* for more information.

Default Settings

Table 13-2 lists the default settings for trunking parameters.

Table 13-2 Default Trunk Configuration Parameters

Parameters	Default
Switch port trunk mode	On
Allowed VSAN list	1 to 4093 user-defined VSAN IDs
Trunking protocol	Enabled

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. Specifically, a PortChannel has the following functionality:

- Provides a point-to-point connection over an 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.

Cisco MDS 9000 Family of switches support 128 PortChannels with 16 interfaces per PortChannel.

This chapter discusses the PortChannel feature provided in the switch.

This chapter contains the following topics:

- PortChannel Examples, page 14-1
- About PortChanneling and Trunking, page 14-3
- About Load Balancing, page 14-5
- Considerations for PortChannel Configurations, page 14-6
- Default Settings, page 14-7

PortChannel Examples

PortChannels on Cisco MDS 9000 Family switches allow flexibility in configuration. 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 14-1 shows portchannel flexibility.

Figure 14-1 PortChannel Flexibility

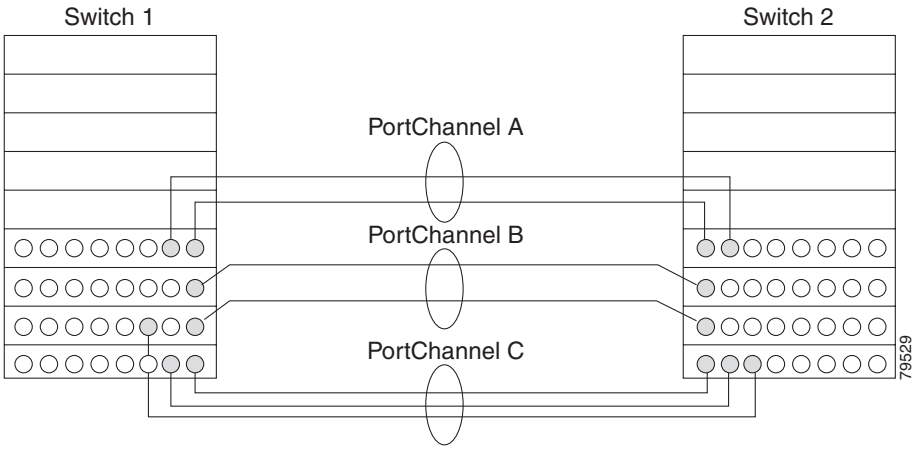

Configuring 32-port Switching Modules and Host-Optimized Ports

The 32-port 1/2-Gbps switching module contains 8 port groups of 4 ports each. When configuring these modules or the host-optimized ports in the Cisco 9100 Series, the following 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 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 as a PortChannel. The other three ports continue to remain in a no shutdown state.

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.

Managing Physical Attributes for a Port

To configure beacon mode and monitor physical attributes for ports from the Fabric Manager, choose **Physical Interfaces** from the menu tree and click the **Physical** tab.

To configure beacon mode and monitor physical attributes for ports from the Device Manager, choose the type of port from the Interface menu and click the **Physical** tab.

The Information pane in Fabric Manager displays attributes for multiple switches. The dialog box from Device Manager displays attributes for a single switch.

To enable or disable beacon mode, check the **BeaconMode** check box. When beacon mode is enabled, an interface LED flashes to identify the interface.

Viewing Port Capability Attributes

To monitor port capability attributes, such as buffer-to-buffer credit, hold time, and class of service from the Fabric Manager, choose **Physical Interface** from the menu tree and click the **Capability** tab.

To monitor these attributes from the Device Manager, choose the type of port from the **Interface** menu and click the **Capability** tab.

The Information pane in Fabric Manager displays attributes for multiple switches. The dialog box from Device Manager displays attributes for a single switch.

About PortChanneling and Trunking

PortChanneling enables several links to be combined into one aggregated link.

Trunking enables an ISL to carry (trunk) multiple VSANs. Trunking can only be configured on 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 14-2.)

Figure 14-2 PortChanneling and Trunking

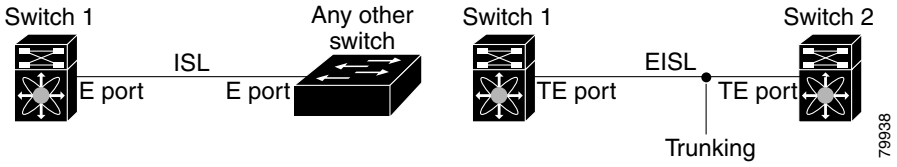

PortChanneling and trunking are used separately across an ISL:

- PortChanneling—Interfaces can be channeled between E ports over multiple ISLs or between TE ports over multiple EISLs.
- Trunking—Trunking, which permits carrying VSAN IDs between switches, can be done only between TE ports over EISLs.

Both PortChanneling and trunking can be used between TE ports over EISLs.

Note When you create a trunked Port Channel in a fabric with FICON VSANs from Fabric Manager, you should immediately go to Device Manager and add the FICON Port Number per that Port Channel. Then, perform a block/unblock, shut/noshut, or Admin Down/Enable to that Port Channel to enable traffic on it.

Managing PortChannel General Attributes

To manage PortChannels, perform the following steps:

Step 1 Choose choose **Switches > PortChannels** from the Fabric Manager menu tree, or choose **Interface > PortChannels** from the Device View. The Information pane in Fabric Manager displays attributes for multiple switches. The dialog box from Device Manager displays attributes for a single switch.

Step 2 To add ports to a PortChannel, click **Create**. You see the Create PortChannel dialog box.

Step 3 To add members to the PortChannel, enter the IP address of the switch into the MemberList field. Identify the other options you want to use and click **OK**.

Managing PortChannel Interface Attributes

To manage PortChannel interface attributes, such as the port mode and trunking from the Fabric Manager, choose **Switches > PortChannels** from the menu tree.

To manage PortChannel interface attributes from the Device Manager, choose **Interface > PortChannels** and click the **Interfaces** tab.

The Information pane in Fabric Manager displays attributes for multiple switches. The dialog box from Device Manager displays attributes for a single switch.

Quiescing/Disabling Port Channel Members

When quiescing, the following apply:

- If an ISL is already quiesced, the message ‘port already quiesced’ is displayed.
- If only one port exists in port channel, a message to that effect is displayed.
- For a port channel between a switch running MD SAN-OS 1.2 and MDS SAN-OS 1.3, the quiesce is performed only on the switch running version 1.3 (or higher). A port shutdown is then run on both switches.

To quiesce a port channel member ISL and administratively bring down both ports, perform these steps:

Step 1 In the Fabric Manager Map pane, right-click on the ISL you want to bring down.

You see the Port Channel menu.

Step 2 Choose **Port Channel > Quiesce and Disable Member**.

A confirmation dialog displays, allowing you to choose the ISL you want to bring down.

Step 3 Choose the ISL you want to bring down.

Step 4 Click **OK** to bring down the ISL, or click **Cancel** to close the dialog without bringing down an ISL.

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 14-3 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 14-3 SID1 and DID1Based Load Balancing

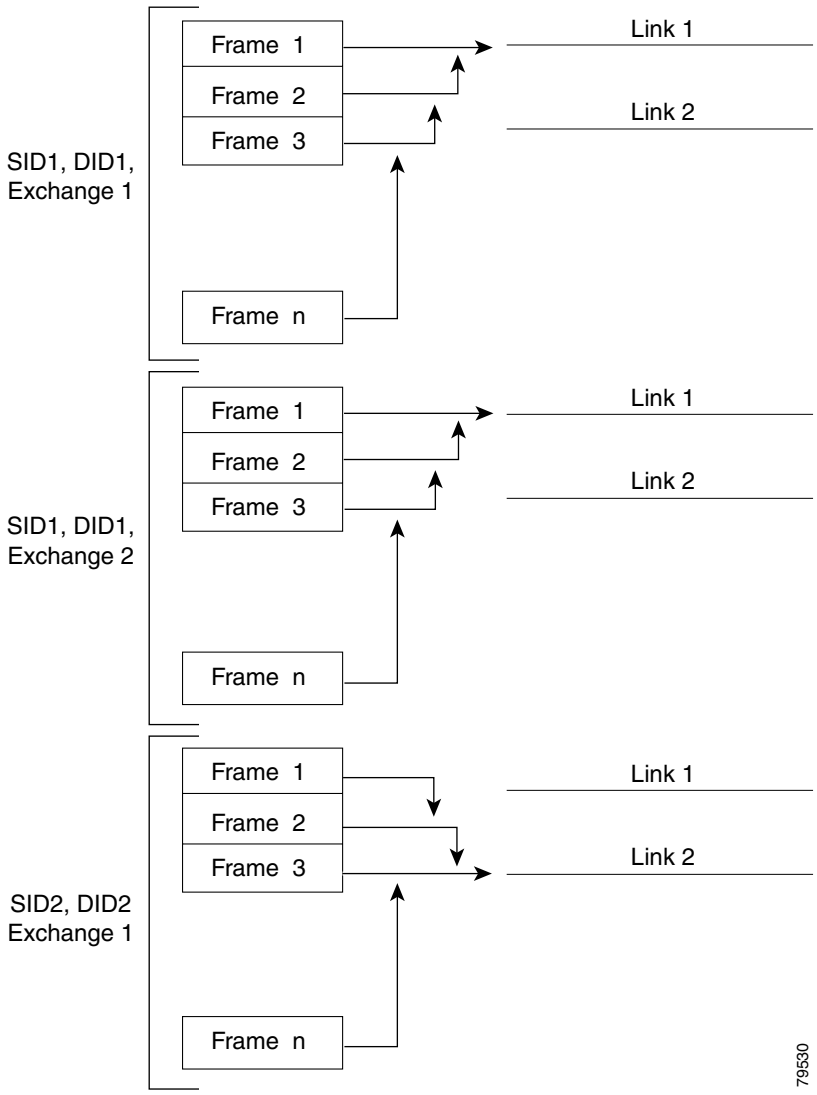

Figure 14-4 illustrates how exchange based load balancing works. When the first frame in an exchange is received for forwarding on an interface, link 1 is chosen by a hash algorithm. All remaining frames in that particular exchange are sent on the same link. For exchange 1, no frame uses link 2. For the next exchange, link 2 is chosen by the hash algorithm. Now all frames in exchange 2 use link 2.

Figure 14-4 SID1, DID1, and Exchange Based Load Balancing

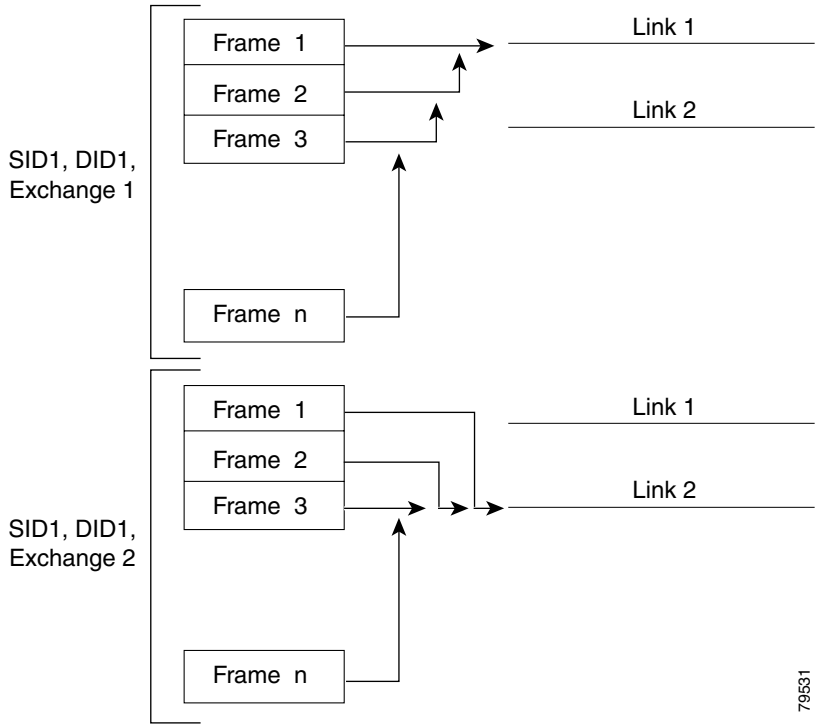

Considerations for PortChannel Configurations

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 two switches. PortChannels require point-to-point connections.

Error Detection

If you misconfigure PortChannels, you may receive the “Error disabled - Possible port channel misconfiguration” message. If you receive this message, the PortChannel’s physical links are disabled since 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.

- Links in a PortChannel must not 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 re-enable the links, and verify that the PortChannel is functioning as required.

If all three conditions are not met, the faulty link is disabled.

Default Settings

Table 14-1 lists the default settings for PortChannels.

Table 14-1 Default PortChannel Parameters

Parameters	Default
PortChannels	FSPF is enabled by default.
Create PortChannel	Administratively up.
Default mode	Auto.
Quiesce	Disabled.

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. This chapter defines various zoning concepts and provides details on zone set and management features in the switch.

This chapter contains the following topics:

- Zoning Features, page 15-2
- Zoning Example, page 15-3
- Configuring a Zone, page 15-4
- Configuring Aliases, page 15-6
- Zone Sets, page 15-7
- Zone Enforcement, page 15-14
- The Default Zone, page 15-14
- Recovering from Link Isolation, page 15-15
- LUN Zoning, page 15-16
- Read-Only Zoning, page 15-17
- Default Settings, page 15-18
- Migrating a Non-MDS Database, page 15-18
- Using the Zone Wizard, page 15-18

Zoning Features

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 the option 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.
- 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.

Zoning Example

Figure 15-1 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 15-1 Fabric with Two Zones

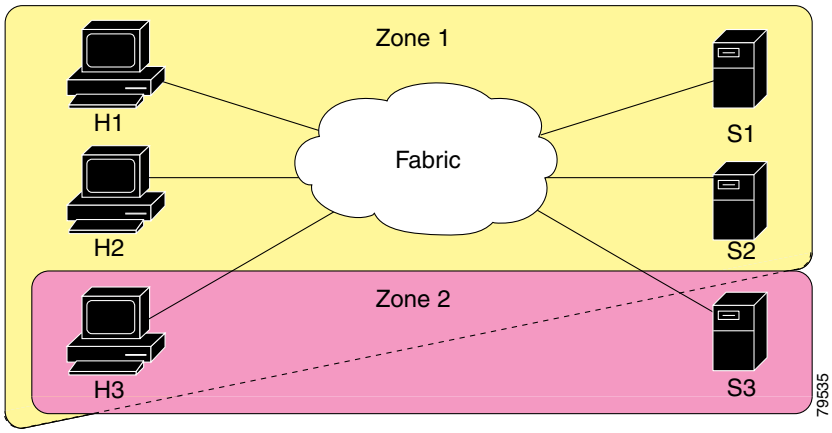

Of course, there are other ways to partition this fabric into zones. Figure 15-2 illustrates another possibility. Assume that there is a need to isolate storage system S2 for the purpose of testing new software. To achieve this, Zone 3 is configured, which contains only host H2 and storage S2. You can restrict access to just H2 and S2 in zone 3, and to H1 and S1 in Zone 1.

Figure 15-2 Fabric with Three Zones

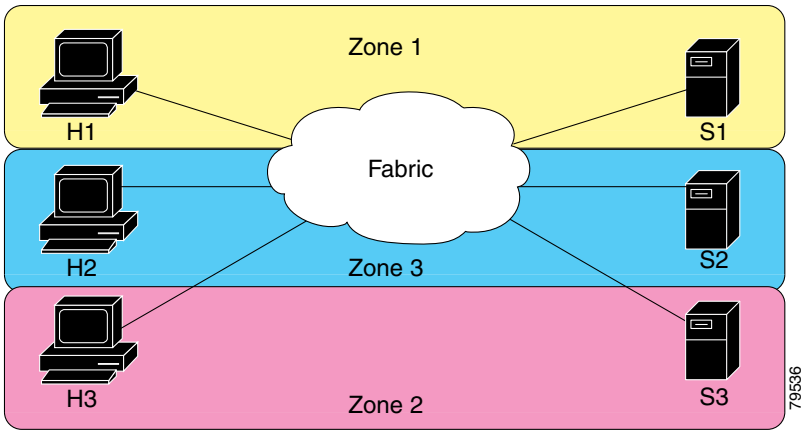

Configuring a Zone

A zone can be configured using one of the following types to assign members:

- pWWN—The WWN of the N or NL port in hex format (for example, 10:00:00:23:45:67:89:ab).
- Fabric port WWN—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).
- 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.
- 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.

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.

If you do not provide a sWWN, the software automatically uses the local sWWN.

Creating Zones

Zones are configured within VSANs, but you can configure zones without configuring any VSANs by configuring them within the default VSAN. The Logical tab displays the VSANs configured in the currently discovered fabric. Note that zone information must always be identical for all the switches in the network fabric.

To create zones, perform the following steps.

Step 1 From the Fabric Manager, choose **Zone > Edit Full Database on Switch**. You see the Select VSAN dialog box. Choose the VSAN and click **OK**.

You can also right-click a VSAN folder in the **Logical** tab and choose **Edit Local Zone Database** from the pop-up menu. You see the Edit VSANxxx Local Full Zones window.

Step 2 Right-click the zone for that VSAN and choose **Insert** to add a zone.

Check the **Set Zone as Read Only** check box to specify that the zone be a ready-only zone. (For more information on read-only zones see the "Read-Only Zoning" section on page 15-17.)

Creating Additional Zones

To create additional zones, follow these steps:

Step 1 With the Edit Full Database on Switch dialog open, right-click the Zones folder and choose **Insert** from the pop-up menu.

Step 2 Enter the zone name in the dialog box that appears and click **OK** to add the zone. The zone is automatically added to the zone database.

Cloning Zones

Another method of adding zones is to clone existing zones.

To clone a zone from the Edit Full Database on Switch window, follow these steps:

Step 1 Click the Zones folder, right-click the folder for the zone that you want to clone, and choose **Clone** from the pop-up menu.

Step 2 Enter the name of the cloned zone. By default, the dialog displays the selected zone as ClonedZone1.

Step 3 Click **OK** to add the cloned zone to the zone database.

Adding Zone Members

Once you have created a zone, you can add members to the zone. You can add members using the following port identification types:

- pWWN—The world wide name of the port configured on the end device (in hex format).
- Fabric port WWN—The world wide name of the physical port on the switch (in hex format).
- FC alias—The alias name in alphabetic characters (for example, Payroll).
- LUN—The logical unit number of a disk in a disk device.

For more information about port identification types, refer to the *Cisco 9000 Family Configuration Guide*.

To add members to a zone, follow these steps:

Step 1 Click the Zones folder, then right-click the folder for the zone to which you want to add members, and choose **Insert** from the pop-up menu.

You see the Add Members to Zone dialog.

Step 2 Check the check box to the left of the NxPort WWN field.

Step 3 Choose one of the ports in the VSAN and click **Add** to add it to the zone. You see the member in the Zone Server database in the lower frame.

Step 4 Repeat these steps to add other members to the zone.

Note When configuring a zone member, you can specify that a single LUN can have multiple IDs depending on the operating system. You can select from 6 different operating systems.

Displaying Port Membership Information

To display port membership information for members assigned to zones, follow these steps.

Step 1 From the Fabric Manager, choose **Zone > Edit Full Database on Switch**. You see the Select VSAN dialog box. Choose the VSAN and click **OK**.

You can also right-click a VSAN folder in the Logical tab and choose **Edit Local Zone Database** from the pop-up menu. You see the Edit VSANxxx Local Full Zones window for the VSAN you selected.

Step 2 Click the **Members** tab.

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. For more information, see the"Changing the Default Zone Policy" section on page 15-15.

Viewing Zone Statistics

To monitor zone statistics from the Zone Server, choose **VSANxxx > Domain Manager** from the Fabric Manager menu tree. You see the zone information in the Information pane. Click on the **Statistics** tab to see the statistics information for the switches in the zone.

Deleting Zones and Members

To delete zones or members, follow these steps.

Step 1 From the Fabric Manager, choose **Zone > Edit Full Database on Switch**. You see the Select VSAN dialog box. Choose the VSAN and click **OK**.

You can also right-click a VSAN folder in the **Logical** tab and choose **Edit Local Zone Database** from the pop-up menu. You see the Edit VSANxxx Local Full Zones window for the VSAN you selected.

Step 2 Choose the zone or member you want to delete.

Step 3 Right-click the object and choose **Delete** from the pop-up menu. The selected object is deleted from the zone database.

Configuring Aliases

You can assign an alias name and configure an alias member using either the FC ID, fabric port WWN (fWWN), or pWWN values.

Creating Zones with Aliases

To create a zone with aliases, perform these steps:

Step 1 Choose **Zone > Edit Local Full Zone Database**. You see the Select VSAN dialog box

Step 2 Choose the VSAN on which you want to create the zone, and click **OK**. You see zone information for that VSAN.

Step 3 Click the **Aliases** tab above the right pane.

Step 4 Right-click in left window pane and choose **Insert**. You see the Create Alias dialog box.

Step 5 Enter the Alias name and choose the pWWN. You can add/associate multiple pWWNs, fWWNs to same alias name. The pWWNs do not have to be attached to the fabric you are currently managing.

Step 6 Click **OK** to create the alias.

Step 7 Right-click on the Zones folder in the left pane and choose **Insert**.

Step 8 Name the zone as desired.

Step 9 Click the **Aliases** tab above the right window pane.

Step 10 Click and drag the desired alias members you created in Steps 5 and 6 above, from the right window pane to the Zone folder you just created in the left window pane.

Step 11 Add the zone to a zone set and activate it accordingly.

Viewing Aliases

Aliases are assigned per port.

To view zone aliases, follow these steps:

Step 1 From the Fabric Manager, choose **Zone > Edit Full Database on Switch**. You see the Select VSAN dialog box. Choose the VSAN and click **OK**.

You can also right-click a VSAN folder in the **Logical** tab and choose **Edit Local Zone Database** from the pop-up menu. You see the Edit VSANxxx Local Full Zones window for the VSAN you selected.

Step 2 Click the **Aliases** tab to see the aliases for that zone.

Zone Sets

In Figure 15-3, two separate sets are created, each with its own membership hierarchy and zone members.

Figure 15-3 Hierarchy of Zone Sets, Zones, and Zone Members

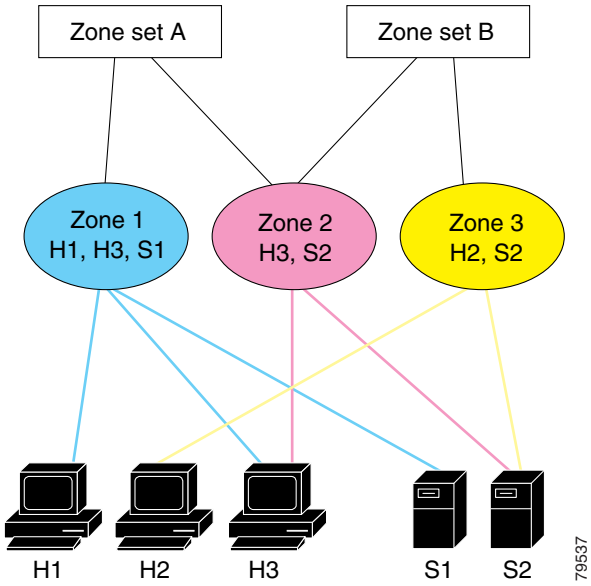

Zones provide a mechanism for specifying access control, while zone sets are a grouping of zones to enforce access control in the fabric. Either zone set A or zone set B can be activated (but not together).

Zone sets are configured with the names of the member zones. If the zone set is in a configured VSAN, the VSAN is also specified.

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. The changes do not take effect until the zone set is activated.
- 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.

Figure 15-4 shows a zone being added to an activated zone set.

If one zone set is active and you activate another zone set, the currently active zone set is automatically deactivated. You don't need to explicitly deactivate the currently active zone set before activating a new zone set.

Figure 15-4 Active and Full Zone Sets

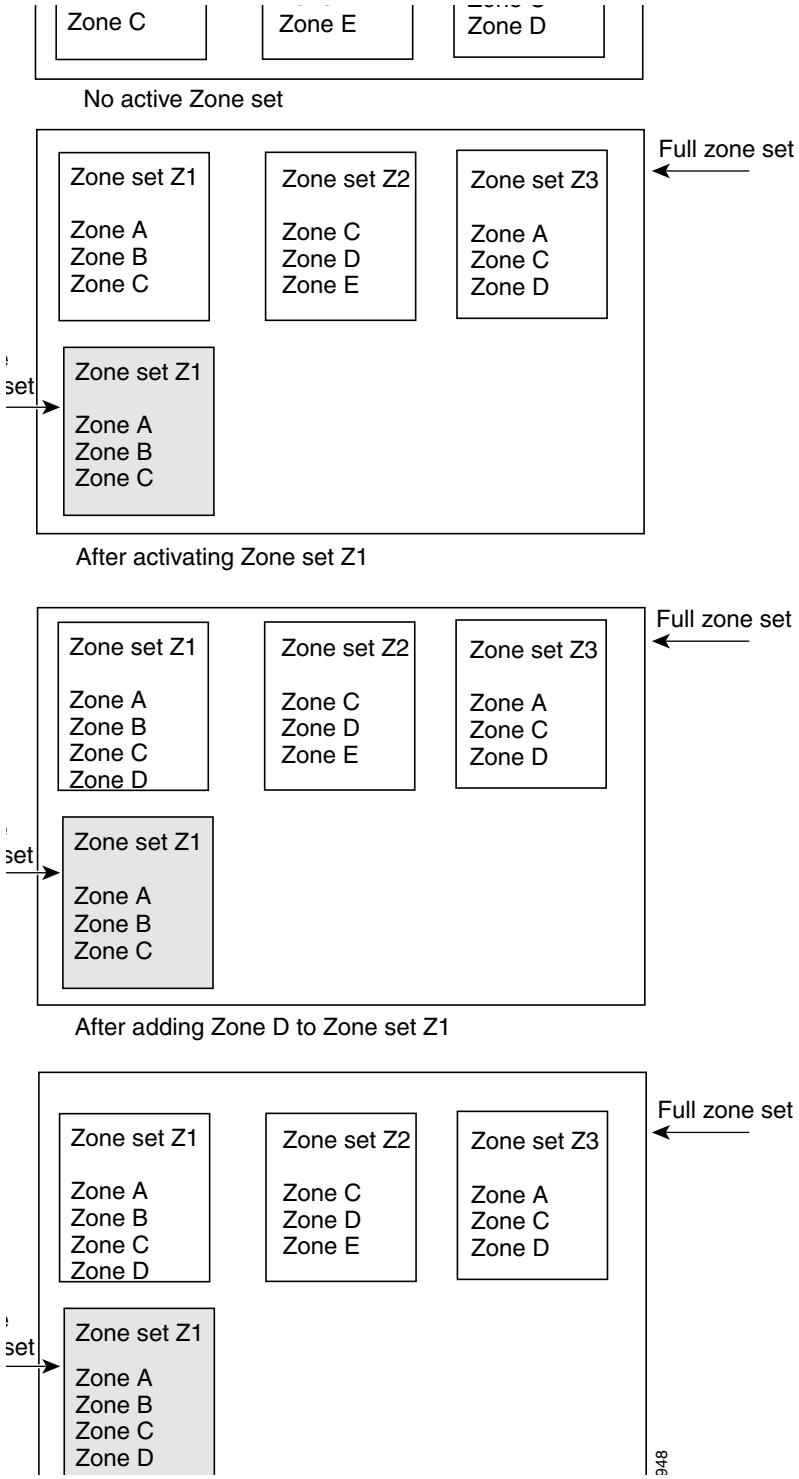

Distributing Zone Sets

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 distribution takes effect while sending merge requests to the adjacent switch or while activating a zone set.

Copying Zone Sets

The active zone set is not a part of the full zone set. You can not make changes to an existing zone set and activate it, if the full zone set is lost or is not propagated. You can make a copy of an active zone set and then edit it without altering the existing active zone set. You can copy an active-zone set to a location in bootflash, volatile, slot0, to a remote location (using FTP, SCP, SFTP, or TFTP), or to the full zone set.

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.

Creating Zone Sets

To create zone sets, perform the following steps.

Step 1 From the Fabric Manager, choose **Zone > Edit Full Database on Switch**. You see the Select VSAN dialog box. Choose the VSAN and click **OK**.

You can also right-click a VSAN folder in the **Logical** tab and choose **Edit Local Zone Database** from the pop-up menu. You see the Edit VSANxxx Local Full Zones window.

Step 2 Right-click the zone set for that VSAN and choose **Insert** to add a zone set.

You can activate the zone set by clicking **Activate**. This configuration is distributed to the other switches in the network fabric.

Note When you confirm the activate operation, the current running configuration is saved to the startup configuration. This permanently saves any changes made to the running configuration (not just zoning changes).

Creating Additional Zone Sets

To create additional zone sets, follow these steps:

Step 1 To create a zone set, right-click the ZoneSets folder in the Edit Full Database on Switch dialog box, and choose **Insert**.

Step 2 Enter the zone set name in the dialog box that appears and click **OK** to add the zone set. The zone set is automatically added to the zone database.

Cloning Zone Sets

Another method of adding zone sets is to clone existing zone sets.

To clone a zone set from the Edit Full Database on Switch window, follow these steps:

Step 1 Click the ZoneSets folder, right-click the folder for the zone set that you want to clone, and choose **Clone** from the pop-up menu.

Step 2 Enter the name of the cloned zone set. By default, the dialog displays the selected zone as ClonedZoneset1.

Step 3 Click **OK** to add the cloned zone set to the zone database.

Adding Zones to a Zone Set

To add a zone to a zone set from the Edit Full Database on Switch window, drag and drop the zone to the folder for the zone set.

Alternatively, follow these steps:

Step 1 Click the ZoneSets folder and then right-click the folder for the zone set to which you want to add a zone and choose **Insert** from the pop-up menu. You see the Zone Server Select Zone dialog box.

Step 2 Select the zone that you want to add to the zone set and click **Add**. The zone is added to the zone set in the zone database.

Activating or Enforcing Zone Sets

Once zones and zone sets have been created and populated with members, you must activate or enforce the zone set. Note that only one zone set can be activated at any time. If zoning is activated, any member that is not assigned to an active zone belongs to the default zone. If zoning is not activated, all members belong to the default zone.

To activate a zone set, follow these steps:

Step 1 Right-click the zone set in the Edit Full Database on Switch dialog box.

Step 2 Click **Activate**. You see the zone set in the Active Zone Set folder.

Note If one zone set is active and you activate another zone set, the currently active zone set is automatically deactivated.

Deactivating Zone Sets

To activate a zone set, follow these steps:

Step 1 Right-click the zone set in the Edit Full Database on Switch dialog box.

Step 2 Click **Deactivate**. The zone set is removed from the Active Zone Set folder.

Importing Active Zone Sets

You can import active zone sets (do a Merge Fail Recovery) if the cause of an ISL failure is a zone merge fail. To import an active zone set, follow these steps:

Step 1 From the Fabric Manager, choose **Zone > Merge Fail Recovery**. You see the Zone Merge Failure Recovery dialog box.

Step 2 Click the **Import Active Zoneset** radio button.

Step 3 Choose the switch from which to import the zone set information from the drop-down list.

Step 4 Choose the VSAN from which to import the zone set information from the drop-down list.

Step 5 Choose the interface to use for the import process.

Step 6 Click **OK** to import the active zone set, or click **Close** to close the dialog without importing the active zone set.

Exporting Active Zone Sets

You can export active zone sets (do a Merge Fail Recovery) if the cause of an ISL failure is a zone merge fail. To export an active zone set, follow these steps:

Step 1 From the Fabric Manager, choose **Zone > Merge Fail Recovery**. You see the Zone Merge Failure Recovery dialog box.

Step 2 Click the **Export Active Zoneset** radio button.

Step 3 Choose the switch to which to export the zone set information from the drop-down list.

Step 4 Choose the VSAN to which to export the Zoneset information from the drop-down list.

Step 5 Choose the interface to use for the export process.

Step 6 Click **OK** to export the active zone set, or click **Close** to close the dialog without exporting the active zone set.

Deleting Zone Sets or Members

To delete zone sets or members, follow these steps.

Step 1 From the Fabric Manager, choose **Zone > Edit Full Database on Switch**. You see the Select VSAN dialog box. Choose the VSAN and click **OK**.

You can also right-click a VSAN folder in the **Logical** tab and choose **Edit Local Zone Database** from the pop-up menu. You see the Edit VSANxxx Local Full Zones window for the VSAN you selected.

Step 2 Choose the zone set or member you want to delete.

Step 3 Right-click the object and choose **Delete** from the pop-up menu. The selected object is deleted from the zone database.

Clearing the Zone Database

Clearing a zone set only erases the full zone database, not the active zone database.

Recovering a Full Zone Database

You can recover a database by copying the active zone database or the full zone database. To recover a zone database, follow these steps:

Step 1 From the Fabric Manager, choose **Zone > Recover Full Zone Database**. You see the Recover Full Zone Database dialog box.

Step 2 Click the **Copy Active** or the **Copy Full** radio button, depending on which type of database you want to copy.

Step 3 Choose the source VSAN from which to copy the information from the drop-down list.

Step 4 If you selected Copy Full, choose the source switch and the destination VSAN from those drop-down lists.

Step 5 Choose the destination switch from the drop-down list.

Step 6 Click **Copy** to copy the database, or click **Close** to close the dialog without copying.

Performing Zone Merge Analysis

To perform a zone merge analysis, follow these steps:

Step 1 From the Fabric Manager, choose **Zone > Merge Analysis**. You see the Zone Merge Analysis window.

Step 2 Choose the first switch to be analyzed from the Check Switch 1 drop-down list.

Step 3 Choose 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 field.

Step 5 Click **Analyze** to analyze the zone merge. Click **Clear** to clear the analysis data from the Zone Merge Analysis window. If you click **Analyze** without clicking **Clear**, the new zone merge analysis data displays below the old data.

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.

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.

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.

Unlike configured zones, default zone information is not distributed to the other switches in the fabric.

Traffic can either be permitted or denied amongst members of the default zone. This information is not distributed to all switches; it must be configured in each switch.

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. 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.

Setting Default Zone Policy

Each VSAN contains a default zone, which by default, contains all connected devices assigned to the VSAN.

You can change the default zone policy for any VSAN by choosing **VSANxxx > Default Zone** from the Fabric Manager menu tree and clicking the **Policies** tab. However, we recommend that you establish connectivity among devices by assigning them to a nondefault 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. The tooltip for each zone indicates the activation time or modification time.

Changing the Default Zone Policy

Each member in the fabric can belong to any zone. If a member does not belong to any zone, it is part of the default zone. If no zone has been activated in the fabric, all members belong to the default zone. Even though a member can belong to multiple zones, a member in the default zone cannot be part of any other zone.

Traffic can be permitted and denied to members in the default zone. This information is not distributed to all switches. Permission and denial must be set for each switch in the fabric.

To permit or deny traffic to members in the default zone from the Zone Server, follow these steps:

Step 1 Choose **VSANxxx > Default Zone** from the Fabric Manager menu tree, and click the **Policies** tab. The zone information displays in the Information pane.

Step 2 Click the DefaultZoneBehavior field and choose either **permit** or **deny** from the pull-down menu.

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. (See Figure 15-5.) 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 "Importing Active Zone Sets" section on page 15-12.
- Export the current database to the neighboring switch. See "Exporting Active Zone Sets" section on page 15-12.
- Manually resolve the conflict by editing the full zone set, activating the corrected zone set, and then bringing up the link.

Figure 15-5 Importing and Exporting the Database

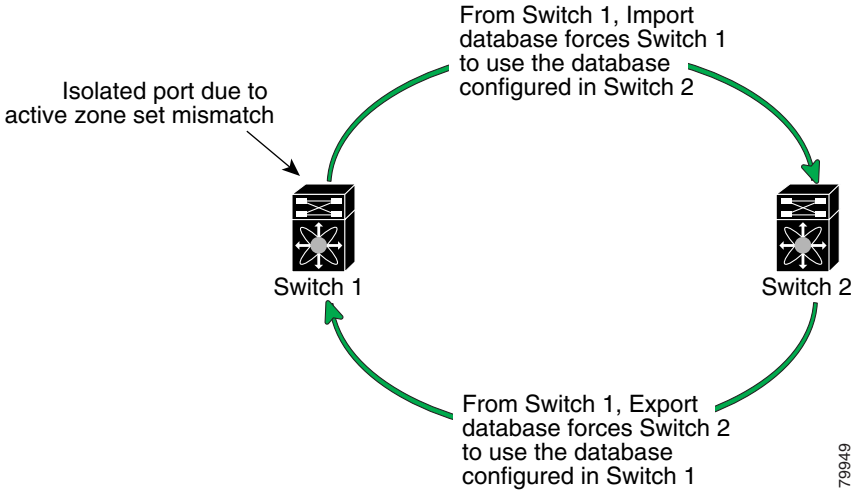

Importing from one switch and exporting from another switch can lead to isolation again.

LUN Zoning

Logical unit number (LUN) zoning is a feature specific to switches in the Cisco MDS 9000 Family.

LUN zoning can be implemented in Cisco MDS 9000 Family switches running Cisco MDS SAN-OS Release 1.2(x) or above.

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.

Figure 15-6 shows a LUN-based zone example.

- Host H1 can access LUN 2 in S1 and LUN 0 in S2. It cannot access any other LUN 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 LUN in S1 or S2.

Unzoned LUNs automatically become members of the default zone.

Figure 15-6 LUN Zoning Access

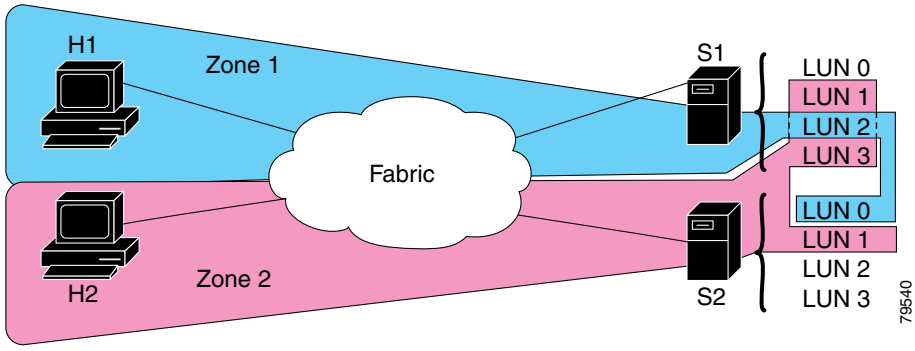

When LUN 0 is not included within a zone, then, as per standards requirements, control traffic to LUN 0 (for example, REPORT_LUNS, INQUIRY) is supported, but data traffic to LUN 0 (for example, READ, WRITE) is denied.

Caution 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.

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 preceding section.

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.

Read-Only Zoning

Read-only zoning can be implemented in Cisco MDS 9000 Family switches running Cisco MDS SAN-OS Release 1.2(x) or above.

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.

Guidelines to Configure 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 above-mentioned Windows operating systems.

Default Settings

Table 15-1 lists the default settings for zone parameters.

Table 15-1 Default Zone Parameters

Parameters	Default
Default zone policy	Denied to all members.
Full zone set distribute	The full zone set(s) is not distributed.
Read-only zones	Read-write attributes for all zones.

Migrating a Non-MDS Database

You use the Zone Migration Wizard to migrate a non-MDS database.

Step 1 From the Fabric Manager, choose **Zone > Migrate Non-MDS Database**. You see the Zone Migration Wizard.

Step 2 Follow the prompts in the wizard to migrate the database.

Using the Zone Wizard

Use the Zone Wizard to configure zones, read-only zones, and IVR zones.

Step 1 From the Fabric Manager, click the **Zone Wizard** icon in the Fabric Manager Zone toolbar (see Figure 15-7).

Figure 15-7 Zone Wizard Icon

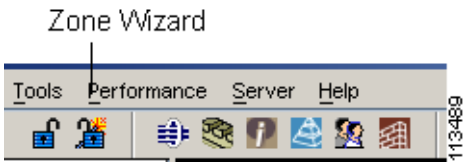

You see the Zone Wizard.

Step 2 Follow the prompts in the wizard to migrate the database.

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 contains the following topics:

- About IVR, page 16-1
- IVR Features, page 16-2
- IVR Terminology, page 16-2
- IVR Guidelines, page 16-3
- Configuring IVR, page 16-4
- Configuring an IVR Topology, page 16-4
- Creating IVZs and IVZSs, page 16-7
- Using the Zone Wizard, page 16-10

About IVR

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.

Data traffic is transported between specific initiators and targets on different VSANs without merging VSANs into a single logical fabric. FC control traffic does not flow between VSANs, nor can initiators access any resource across VSANs aside from the designated ones. Valuable resources like tape libraries are easily shared across VSANs without compromise.

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.

IVR Features

IVR has the following features:

- Accesses resources across VSANs without compromising other VSAN benefits.
- Transports data traffic between specific initiators and targets on different VSANs without merging VSANs into a single logical fabric.
- Shares valuable resources (like tape libraries) across VSANs without compromise.
- Provides efficient business continuity or disaster recovery solutions when used in conjunction with FCIP.
- Compliant with Fibre Channel standards compliant.
- Incorporates with third-party switches--if the IVR-enabled VSANs are configured in **interop 1** mode.

IVR Terminology

The terms used in this chapter are explained in this section.

- Native VSAN—The VSAN to which an end device logs on is called a native VSAN for that end device.
- Inter-VSAN zone (IVZ)—Defines a set of end devices s 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 association. You can configure up to 200 IVZs and 2000 IVZ members on any switch in the Cisco MDS 9000 Family.
- Inter-VSAN zone sets (IVZS)—One or more IVZs make up an IVZS. You can configure up to 32 IVZSs on any switch in the Cisco MDS 9000 Family. Only one IVZS can be active at any time.
- IVR path—An IVR path is a set of switches and inter-switch links via 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—An edge VSAN refers to a VSAN which 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 16-1, VSANs 1, 2, and 3 are edge VSANs.

  An edge VSAN for one IVR path can be a transit VSAN for another IVR path.
- Transit VSAN—Transit VSAN is a VSAN that exists along an IVR path from the source edge VSAN of that path to the destination edge VSAN of that path.

  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.
- 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.

IVR Guidelines

Before configuring an IVR SAN fabric, consider the following guidelines:

- Configure unique domain IDs across all VSANs and switches participating in IVR operations. 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.
- An Enterprise License Package is required for this feature.

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.

IVR-enabled VSANs must be configured in **no interop** (default) mode or **interop 1** mode.

Domain ID Guidelines

Domain IDs must be unique across inter-connected VSANs. To ensure unique domain IDs across inter-connected VSAN, follow 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.

Transit VSANs Guidelines

Consider the following guidelines for transit VSANs:

- Besides defining the IVZ membership, you can choose to specify a set of transit VSANs to provide connectivity between two edge VSANs:
  - If two edge VSANs in an IVZ overlap, then a transit VSAN is not required (though, not prohibited) to provide connectivity.
  - If two edge VSANs in an IVZ do not overlap, you may need one or more transit VSANs to provide connectivity. Two edge VSANs in an IVZ will not overlap if IVR is not enabled on a switch that is a member of both the source and destinations edge VSANs.
- Traffic between the edge VSANs only traverses through the shortest IVR path.
- Transit VSAN information is common to all IVZs. Sometimes, a transit VSAN can also double-up as an edge VSAN in another IVZ.

Border Switch Guidelines

Before configuring border switches, consider the following guidelines:

- Border switches require SAN-OS Release 1.3(1) or higher.
- A border switch must be a member of two or more VSANs.
- A border switch that facilities IVR communications must be IVR enabled.

- To provide redundant paths between active IVZ members, IVR can (optionally) be enabled on additional border switches.
- The VSAN topology configuration must be updated before a border switch is added or removed.

Configuring IVR

To configure IVR in a SAN fabric, follow these steps.

Step 1 Verify that unique domain IDs are configured in all switches and VSANs participating in IVR.

Step 2 Enable IVR in the border switches.

Step 3 Create and activate the required IVR topology in all the IVR-enabled border switches.

Step 4 Create and activate IVZSs in all the IVR-enabled border switches.

Step 5 Verify the IVR configuration.

Unique Domain ID Configuration Options

You can configure domain IDs using one of two options:

- Configure allowed-domains list using the Domain Manager MIBs so that the domains in different VSANs are non-overlapping on all participating switches and VSANs.
- Configure static, non-overlapping domains (using the CLI) for each participating switch and VSAN.

Enabling IVR

The IVR feature must be enabled in all border switches in the fabric that participate in the IVR. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family. To begin configuring the IVR feature, you must explicitly enable IVR on the required switches in the fabric.

The configuration and verification commands for the IVR feature are only available when IVR is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.

Configuring an IVR Topology

This section explains the process used to create an IVR topology.

Creating an IVR Topology

You must create the IVR topology in every IVR-enabled switch in the fabric. You can have up to 64 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 autonomous fabric ID (AF ID) which distinguishes two VSANs that are logically and physically separate, but have the same VSAN number. SAN-OS Release 1.3(1) supports only one AF ID.

The use of a single AF ID does not allow for segmented VSANs in an inter-VSAN topology.

Ensure to repeat this configuration in all IVR-enabled switches.

Transit VSANs are deduced based on your configuration The IVR feature does not have an explicit transit-VSAN configuration. In the example used above, VSAN 2 is the transit VSAN between VSANs 1 and 3.

Creating IVR Zones and Zone Sets

To create IVR zones or zone sets, perform the following steps:

Step 1 From the Fabric Manager, choose **Zone > IVR (Inter VSAN Routing) > Edit Full Database on Switch**.

The Edit VSANxxx Local Full Zones window displays for the VSAN you selected.

Step 2 Right-click the zone set or zone for that VSAN and choose **Insert** to add a zone set or zone.

If you are adding a zone set, you can activate it by clicking the **Activate** button. This configuration is distributed to the other switches in the network fabric.

Note When you confirm the activate operation, the current running configuration is saved to the startup configuration. This permanently saves any changes made to the running configuration (not just zoning changes).

Note Some time zones beginning with prefix 'IVRZ' and a zone set with name 'nozoneset' appear in logical view. The zones with prefix 'IVRZ' are IVR zones which get appended to regular active zones. The prefix 'IVRZ' is appended to active IVR zones by the system. Similarly the zone set with name 'nozonset' is an IVR active zone set created by system if no active zone set is available for that VSAN and if 'ivrZoneSetActivateForce' flag is enabled on switch. In 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.

Note Do not create a zone with prefix 'IVRZ' or a zone set with name 'nozonset'. These names are used by the system for identifying IVR zones.

Creating Additional IVR Zones and Zone Sets

To create additional zones and zone sets, follow these steps:

Step 1 With the Zone > IVR (Inter VSAN Routing) > Edit Full Database on Switch dialog open, right-click the Zones folder and choose **Insert** from the pop-up menu.

Step 2 Enter the zone name in the dialog box that appears and click **OK** to add the zone.

The zone is automatically added to the zone database.

Step 3 To create a zone set, right-click the ZoneSets folder in the Edit Full Database on Switch dialog, and choose **Insert**.

Step 4 Enter the zone set name in the dialog box that appears and click **OK** to add the zone set.

The zone set is automatically added to the zone database.

Activating IVR Zone Sets

Once the zone sets have been created and populated, you must activate the zone set.

To activate an IVR zone set, follow these steps:

Step 1 Right-click the zone set in the Zone > IVR (Inter VSAN Routing) > Edit Full Database on Switch dialog.

Step 2 Click **Activate**.

Note The active zone set in Edit Zone is always shown in bold, even after successful activation. This is because a member of this VSAN must be participating in IVR zoning. Since the IVR zones get added to active zones, the active zone set configuration is always different from local zone set configuration with same name.

Deactivating IVR Zone Sets

To activate a zone set, follow these steps:

Step 1 Right-click the zone set in the Zone > IVR (Inter VSAN Routing) > Edit Full Database on Switch dialog.

Step 2 Click **Deactivate**.

Recovering an IVR Full Zone Database

You can recover an IVR zone database by copying the IVR full zone database.

To recover an IVR zone database, perform these steps.

Step 1 From the Fabric Manager, choose **Zone > IVR (Inter VSAN Routing) > Recover Full Zone Database**. You see the Recover Full Zone Database dialog.

Step 2 Click the **Copy Active** or the **Copy Full** radio button, depending on which type of database you want to copy.

Step 3 Choose the source VSAN from which to copy the information from the drop-down list.

Step 4 If you selected Copy Full, choose the source switch and the destination VSAN from those drop-down lists.

Step 5 Choose the destination switch from the drop-down list.

Step 6 Click **Copy** to copy the database, or click **Close** to close the dialog without copying.

Recovering an IVR Full Topology

You can recover a topology by copying the active zone database or the full zone database.

To recover a zone database, perform these steps.

Step 1 From the Fabric Manager, choose **Zone > IVR (Inter VSAN Routing) > Recover FullTopology**. You see the Recover Full Topology dialog box.

Step 2 Click the **Copy Full** radio button.

Step 3 Choose the source VSAN from which to copy the information from the drop-down list.

Step 4 Choose the source switch and the destination VSAN from those drop-down lists.

Step 5 Choose the destination switch from the drop-down list.

Step 6 Click **Copy** to copy the topology, or click **Close** to close the dialog without copying.

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 IVZS 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 the **interop-mode 1** option is enabled.

IVR Using LUN Zoning or Read-Only Zoning

LUN-zoning and read-only zoning can be used between members of active IVR zones. To configure this service, you need to create and activate LUN-zones and/or read-only zones between the desired IVZ members in all relevant edge VSANs using the zoning interface.

The LUN zoning and read-only zoning features cannot be configured in a IVZS setup.

Creating IVZs and IVZSs

As part of the IVR configuration, you need to configure one or more IVZs to enable cross-VSAN communication. To achieve this result, you must specify each IVZ as a set of (pWWN, VSAN) entries. Like zones, several IVZs can be configured to belong to an IVRS. You can define several IVZSs and activate only one of the defined IVZSs.

The same IVZS must be activated on all the IVR-enabled switches.

Zones versus IVZs

Table 16-1 identifies the key differences between IVZs and Zones.

Table 16-1 Key Differences between IVZs and Zones

IVZs	Zones
IVZ membership is specified using the VSAN and pWWN combination.	Zone membership is specified using pWWN, fabric WWN, sWWN, or the fabric ID.
Default zone policy is always deny (not configurable).	Default zone policy is deny (configurable).

Automatic IVZ Creation

Figure 16-1 depicts an IVZ 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.

Figure 16-1 IVZ

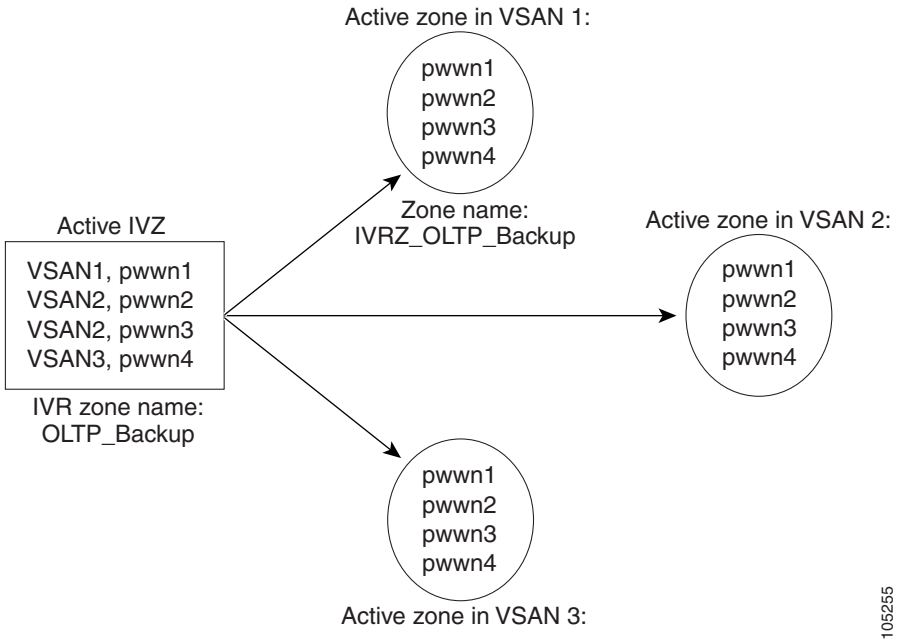

A zone corresponding to each active IVZ is automatically created in each edge VSAN specified in the active IVZ. All pWWNs in the IVZ are members of these zones in each VSAN.

The zones are created automatically by the IVR process when an IVZS 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 IVZS configuration when a new zone set is activated. Like zone sets, IVR zone sets are also activated non-disruptively.

If pwwn1 and pwwn2 are in an IVZ in the current as well as the new IVZS then activation of the new IVZS does not cause any traffic disruption between them.

Configuring and Activating IVZs and IVZSs

IVZ and IVZS names are restricted to 64 alphanumeric characters.

Using the force Option

Use the **force** option to activate the specified IVZS. Table 16-2 lists the various scenarios with and without the force option.

Table 16-2 Using the force Option

Case	Default Zone Policy	Active Zone Set before IVRZ Activation	Force Option Used?	IVZS Activation Status	Active IVRZ Created?	Possible Traffic Disruption
1	Deny	No active zone set	No	Failure	No	No
2	Yes	Success	Yes	No		
3	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		

Using the **force** option of IVZS activation 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 IVZS activation will fail. However, IVZS activation will go through if the **force** option is used. Since zones are created in the edge VSANs corresponding to each IVZ, traffic may be disrupted in edge VSANs where the default zone policy is permit.

Be sure to repeat this configuration in all border switches participating in the IVR configuration.

Using the Cisco MDS Fabric Manager, you can distribute IVZ configurations to all IVR-capable switches in the interconnected VSAN network. Refer to the *Cisco MDS 9000 Family Fabric Manager User Guide* for more information.

Clearing the IVZ Database

Clearing a zone set only erases the configured zone database, not the active zone database.

Using the Zone Wizard

Use the Zone Wizard to configure zones, read-only zones, and IVR zones.

Step 1 From the Fabric Manager, click the **Zone Wizard** icon in the Fabric Manager Zone toolbar.

Figure 16-2 Zone Wizard icon

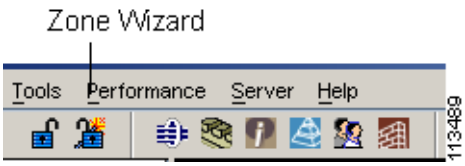

The Zone Wizard displays.

Step 2 Follow the prompts in the wizard to migrate the database.

Managing FLOGI, Name Server, FDMI, and RSCN Databases

This chapter describes the fabric login database, the name server features, and Registered State Change Notification (RSCN) information provided in the Cisco MDS 9000 Family.

This chapter contains the following topics:

- Displaying FLOGI Details, page 17-1
- Configuring the Name Server Proxy Feature, page 17-2
- Displaying FDMI, page 17-2
- Displaying RSCN Information, page 17-2
- Viewing General Attributes for the Name Server, page 17-3
- Viewing Advanced Attributes for the Name Server, page 17-3
- Proxy Ports for the Name Server, page 17-3
- Viewing Name Server Statistics, page 17-4
- Viewing RSCN Nx Registrations, page 17-4
- Viewing RSCN Statistics, page 17-4
- Viewing FLOGI Attributes, page 17-4
- Viewing Port ELP Attributes, page 17-5
- Viewing Trunk Configuration, page 17-5

Displaying FLOGI Details

In a Fibre Channel fabric, each host or disk requires an FC ID. Choose **Switches > FC Logical > FLOGI** to verify if a storage device is displayed in the Fabric login (FLOGI) table as in the following examples. 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.

Configuring the Name Server Proxy Feature

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 which 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.

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.

Displaying 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.

Displaying FDMI

SAN-OS 1.3(x) provides support for the Fabric-Device Management Interface (FDMI) functionally, as described in the FC-GS-4 standard. FDMI enables management of devices such as Fibre Channel Host Bus Adaptors (HBAs) through inband 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 RSCN Information

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

- Or any other similar event that affects the operation of the host

Apart from sending these events to registered hosts a switch RSCN (SW-RSCN) is sent to all reachable switches in the fabric.

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.

The SCR table cannot be configured, it is only populated if one or more hosts send SCR frames to register for RSCN information. If there are no entries in the SCR tables, no host is interested in receiving RSCN information.

Sending RSCNs

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, you have two disks (D1, D2) and a host (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:

- If the multi-pid option is disabled on switch 1, then two RSCNs is generated to Host H--one for the disk D1 and another for disk D2.
- If the multi-pid option is enabled on switch 1, then a single RSCN is generated to host H, and the RSCN payload lists the affected port IDs (in this case, both D1 and D2).

Some Nx ports may not understand multi-pid RSCN payloads. If so, you must disable the multi-pid RSCN option.

Viewing General Attributes for the Name Server

To view general name server attributes from the Device Manager, choose **Name Server** from the FC menu. The Name Server dialog box, with the General tab selected, displays name server attributes for a single switch.

Viewing Advanced Attributes for the Name Server

To monitor advanced name server attributes from the Device Manager, choose **Name Server** from the FC menu and click the **Advanced** tab. The Name Server dialog box, with the Advanced tab selected, displays advanced name server attributes for a single switch.

Proxy Ports for the Name Server

To configure proxy ports for the name server from Fabric Manager, choose **FC > Name Server** on the menu tree and click the **Proxies** tab. The Information pane from the Fabric Manager displays name server proxy ports for multiple switches.

To configure proxy ports for the name server from the Device Manager, choose **Name Server** from the FC menu and click the **Proxy** tab. The Name Server dialog box, with the Proxy tab selected, displays name server proxies for a single switch.

Configure proxy attributes for the name server.

Viewing Name Server Statistics

To monitor name server statistics from the Fabric Manager, choose **FC > Name Server** on the menu tree and click the **Statistics** tab. The Information pane from the Fabric Manager displays name server statistics for multiple switches.

To monitor name server statistics from the Device Manager, choose **Name Server** from the FC menu and click the **Statistics** tab. The Name Server dialog box, with the Statistics tab selected, displays name server statistics for a single switch.

Viewing RSCN Nx Registrations

To view Nx registrations for RSCN from the Fabric Manager, choose **FC > RSCN** on the menu tree, and click the **Registrations** tab. The Information pane from the Fabric Manager displays Nx registrations for RSCN for multiple switches.

To monitor Nx registrations for RSCN from the Device Manager, choose **RSCN** from the FC menu. The RSCN dialog box, with the Nx Registrations tab selected, displays Nx registrations for RSCN for a single switch.

Viewing RSCN Statistics

To monitor registered state change notification (RSCN) statistics from the Fabric Manager, choose **FC > RSCN** on the menu tree and click the **Statistics** tab. The Information pane from the Fabric Manager displays RSCN statistics for multiple switches.

To monitor RSCN from the Device Manager, choose **RSCN** from the FC menu and click the **Statistics** tab. The RSCN dialog box, with the Statistics tab selected, displays RSCN statistics for a single switch.

Viewing FLOGI Attributes

To view fabric login (FLOGI) attributes, such as the Fibre Channel ID (FCID), port name, and class of service for FxPorts from the Fabric Manager, choose **FC > Physical Interfaces** on the menu tree, and click the **FLOGI** tab.

To view FLOGI attributes from the Device Manager, choose **FxPorts** or **All Ports** from the Interface menu and click the **FLOGI** tab.

The Information pane in Fabric Manager displays attributes for multiple switches. The dialog box from Device Manager displays attributes for a single switch.

Viewing Port ELP Attributes

To monitor exchange link parameter (ELP) attributes, such as port and node world wide names and class of service from the Fabric Manager, choose **FC > Physical Interfaces** from the menu tree and click the **ELP** tab.

To monitor these attributes from the Device Manager, choose **xEPorts** or **All Ports** from the Interface menu and click the **ELP** tab.

The Information pane in Fabric Manager displays attributes for multiple switches. The dialog box from Device Manager displays attributes for a single switch.

Viewing Trunk Configuration

To monitor trunking for ports from the Fabric Manager, choose **FC > Physical Interfaces** from the menu tree, and then click the **Trunk Failures** tab.

To view trunking for ports from the Device Manager, choose **xEPorts** from the Interface menu and then click the **Trunk Failures** tab.

The Information pane in Fabric Manager displays attributes for multiple switches. The dialog box from Device Manager displays attributes for a single switch.

Configuring Switch Security

The authentication, authorization, and accounting (AAA) strategy is used to verify identity of, grant access, and track the actions of remote users in all switches in the Cisco MDS 9000 Family. The Remote Access Dial-In User Service (RADIUS) and Terminal Access Controller Access Control System Plus (TACACS+) provide AAA solutions.

Based on the user ID and password combination provided, switches perform local authentication using a local database or remote authentication using AAA server(s). A global, preshared, secret key authenticates communication between the AAA servers. This secret key can be configured for all AAA server groups or for only a specific AAA server. This kind of authentication provides a central configuration management capability.

This chapter contains the following topics:

- Switch Management Security, page 18-2
- Switch AAA Functionalities, page 18-2
- Configuring RADIUS, page 18-4
- Configuring TACACS+, page 18-7
- Configuring Server Groups, page 18-9
- Local AAA, page 18-9
- Authentication and Authorization Process, page 18-9
- Configuring Role-Based CLI Authorization, page 18-11
- Recovering Administrator Password, page 18-12
- Configuring SSH Services, page 18-12
- About SNMP Security, page 18-13
- Default Security Settings, page 18-20
- Restricting Switch Access, page 18-21

Switch Management Security

Management security in any switch in the Cisco MDS 9000 Family is implemented using the Command-line interface (CLI) or Simple Network Management Protocol (SNMP).

SNMP Security

The SNMP agent supports security features for SNMPv1, SNMPv2c, and SNMPv3. Normal SNMP security mechanisms apply to all applications that use SNMP (for example, Cisco MDS 9000 Fabric Manager).

Users and roles configured through the CLI are different from users and roles configured through SNMP. These configurations do not directly correspond with each other. However, you can configure both CLI and SNMP identically, if required.

CLI Security

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 TACACS+), or none.

- Remote security control
  - Using Remote Authentication Dial-In User Services (RADIUS).
  - Using Terminal Access Controller Access Control System plus (TACACS+).
- Local security control.
- Trivial authentication.

These authentication mechanisms can also be used to configure AAA for the following scenarios:

- iSCSI authentication
- Fibre Channel Security Protocol (FC-SP) authentication

Switch AAA Functionalities

Using CLI, you can configure Authentication, Authorization, and Accounting (AAA) switch functionalities on any switch in the Cisco MDS 9000 Family.

This section contains the following topics:

- Authentication, page 18-3
- Authorization, page 18-3
- Accounting, page 18-3
- Remote Authentication by AAA Servers, page 18-3
- Remote Authentication Guidelines, page 18-3
- Server Groups, page 18-4
- AAA Service Configuration Options, page 18-4

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 lookup database) or remote authentication (using one or more RADIUS or TACACS+ servers).

Authorization

By default, two roles exist in all 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.

The two default roles cannot be changed or deleted. You can create additional roles and configure the following options:

- Assign user roles either 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 that user is authenticated through remote AAA server.

Accounting

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

Remote Authentication by AAA Servers

AAA authentication provides the following advantages over local database authentication:

- Requires only one password to be shared between the switch and the AAA servers.
- Easier to manage user password lists for each switch in the fabric.
- AAA servers are deployed widely across enterprises and can be easily adopted.

Remote Authentication Guidelines

When you prefer using remote C servers, follow these guidelines:

- A minimum of one AAA server should be IP reachable.
- If all configured AAA servers are not reachable, the policy configured on the switch determines the authentication method.

- RADIUS servers are easily reachable if an overlay Ethernet LAN is attached to the switch. This is the recommended method.
- SAN networks connected to the switch should have at least one gateway switch connect to the Ethernet LAN containing the AAA servers. If you are using IP connectivity to reach an AAA server, the SAN connects to the switch.

Server Groups

You can specify remote AAA servers for authentication, authorization and accounting using server groups. A server group consists 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 response. 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 fails to respond, then that server group option is considered a failure. You can create a server group using the **aaa group server** command.

If required, you can specify multiple server groups. If the MDS switch encounters errors from the server(s) in the first group, it tries the servers in next server group.

AAA Service Configuration Options

AAA configuration in Cisco MDS switches is service based. You can have separate AAA configurations for following services:

- Telnet or SSH login—Choose **Switches > Security > SSH**.
- iSCSI authentication—Choose **End Devices > iSCSI > Global**.
- FC-SP authentication—Chose **Switches > Security > FC-SP**.

In general, server group, local, and none are the three options that can be specified for any service in an AAA configuration. Each option will be tried in the order specified. If all the methods fail, local is tried

Even if local is not specified as one of the options, it is tried when all other configured options fail.

Configuring RADIUS

Cisco MDS switches use the RADIUS protocol to communicate with remote AAA servers. You can configure multiple RADIUS servers and set timeout and retry counts.

This section defines the RADIUS operation, identifies its network environments, and describes its configuration possibilities.

This section contains the following topics:

- About RADIUS, page 18-5
- Configuring RADIUS Authentication, page 18-5
- Configuring RADIUS Servers, page 18-5
- Setting the RADIUS Server Address, page 18-5
- Setting the RADIUS Preshared Key, page 18-6

- Setting Iterations of the RADIUS Server, page 18-6
- Defining Vendor-Specific Attributes, page 18-6

About RADIUS

RADIUS is a distributed client/server system that secures networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco MDS switches and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.

RADIUS is a fully open protocol, distributed in source code format, that can be modified to work with any security system currently available on the market.

You can set the RADIUS server address, the RADIUS preshared key, the RADIUS server time-out interval, iterations of the RADIUS server, define vendor-specific attributes, and display RADIUS server details.

Configuring RADIUS Authentication

To configure RADIUS authentication from the Fabric Manager, choose **Security > Radius** from the menu tree.

To configure RADIUS authentication from the Device Manager, choose **Security > Radius (CLI)**.

Configuring RADIUS Servers

To configure RADIUS servers, perform the following steps:

Step 1 From the Device Manager, choose **Security > Radius** and click the **Servers** tab. You see the Radius dialog box with the Servers tab selected.

To configure RADIUS servers from the Fabric Manager, choose **Security > Radius** from the menu tree and click the **Servers** tab. You see the Radius information in the Information pane.

Step 2 To add a Radius server, click **Create** on the Device Manager dialog box, or click the **Create Row** icon on the Fabric Manager toolbar.

You see the Create Radius Server dialog box.(In Fabric Manager, you can specify the switches to which the configuration applies.)

Step 3 Complete the fields, and click **OK**.

Setting the RADIUS Server Address

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. From Fabric Manager, choose **Switches > Security > Radius > Servers** to set RADIUS server addresses.

Setting the RADIUS 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 65 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. From Fabric Manager, choose **Switches > Security > Radius** and click the **Defaults** tab to override this global key assignment.

Setting Iterations of the RADIUS Server

By default, a switch retries a RADIUS server connection only once. This number can be configured. The maximum is five retries per server. You can revert the retry number to its default. From Fabric Manager, choose **Switches > Security > Radius > Defaults**.

Defining 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
- *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, like authorization information, along with authentication results. This authorization information is specified through VSAs.

VSA Format

The following VSA protocol options are supported:

- 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:

- 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. This is an example using the roles attribute:

  ```
  Cisco-AVPair = shell:roles=”network-admin vsan-admin”
  ```

- 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 value.

Configuring TACACS+

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.

This section contains the following topics:

- About TACACS+, page 18-7
- Advantages of TACACS+, page 18-7
- Enabling TACACS+, page 18-8
- Setting the TACACS+ Server Address, page 18-8
- Setting the Secret Key, page 18-8
- Setting the Timeout Value, page 18-8
- Defining Custom Attributes for Roles, page 18-8

About TACACS+

TACACS+ is a client-server protocol which 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 SAN-OS 1.3(x) enables the following advantages over RADIUS authentication:

- Provides independent, modular AAA facilities--authorization can be done without authentication.
- Performs independent of servers if it is configured to its own database.
- TCP transport protocol to send data between the AAA client and server, using 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.

Advantages of TACACS+

This section provides a brief list of advantages that TACACS+ has over and RADIUS.

- Uses TCP protocol which has a connection-oriented transport
- Provides built-in transport support
- Provides a separate acknowledgment that a request has been received
- Provides immediate indication of a crashed, or not running, server
- Detects server crashes out-of-band with actual requests

- Maintains simultaneous connections to multiple servers
- Adapts to growing, as well as congested networks

Enabling TACACS+

By default, the TACACS+ feature is disabled in all switches in the Cisco MDS 9000 Family. You must explicitly enable the TACACS+ feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.

Setting the TACACS+ Server Address

If a secret key is not configured for a configured server, a warning message is issued and the global secret encryption key is automatically used.

Setting the Secret Key

From Fabric Manager, choose **Switches > Security > TACACS+ > Defaults** to configure global values for the key for all TACACS+ servers.

Secret keys configured for individual servers override the globally configured values.

Setting the Timeout Value

From Fabric Manager, choose **Switches > Security > TACACS+ > Defaults** to configure global timeout values for all TACACS+ servers.

Timeout values configured for individual servers override the globally configured values.

Defining Custom Attributes for Roles

MDS uses TACACS+ custom attribute for service shell to configure the roles to which a user belongs. TACACS+ attributes are specified as 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"
```

TACACS+ custom attributes can be defined on an ACS server for various services (for example, shell). MDS requires the TACACS+ custom attribute for service shell to be used for defining roles.

Configuring 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.

You can configure these server groups at any time but they only take effect when you apply them to a AAA service. From Fabric Manager, choose **Switches > Security > AAA > Server Groups**.

You can specify one or more remote AAA servers to authenticate users using server groups.

Local AAA

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.

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 trying to manage 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).

The following steps explain the authorization and authentication process. shows a flow chart of the process.

Figure 18-1 Switch Authorization and Authentication Flow

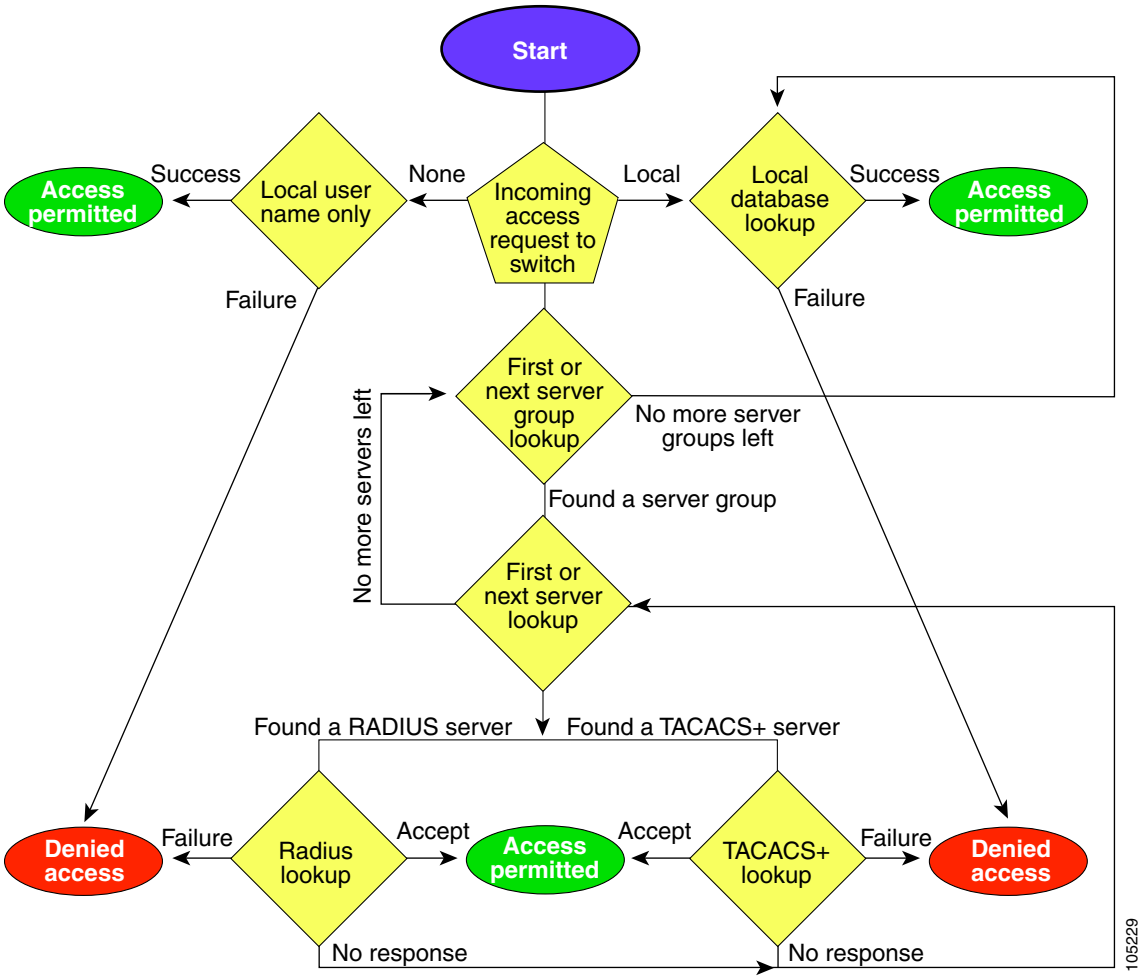

Step 1 When you can log in to the required switch in the Cisco MDS 9000 Family, you have the option to use the Telnet, SSH, or Console login options.

Step 2 When you configure 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 will be 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 fails, then local database is used for authentication.

Step 3 When you are successfully authenticated through a remote AAA server, then the following possibilities apply:

- If AAA server protocol is RADIUS, the user roles specified in cisco-av-pair attribute is downloaded with 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 shell
- If user roles were not retrieved successfully from remote AAA server, then the user will have role of network-operator assigned once he logs in.

Step 4 If your user name and password are successfully authenticated, you are allowed to log in.

Configuring Role-Based CLI 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.

Each role can contain multiple users and each user can be part of multiple roles. For example, if role1 users are only allowed to perform configuration commands, and role2 users are only allowed to perform debug commands, then if Joe belongs to both role1 and role2, he can perform configuration as well as debug commands.

If you belong to multiple roles, you can execute a superset 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 Rules and Features for Each Role

A rule specifies operations that can be performed by a specific role. Each rule consists of a command type (for example, **config**, **clear**, **show**, **exec**, **debug**), and an optional feature name (for example, FSPF, zone, VSAN, fcping, interface).

Configuring the VSAN Policy

Configuring the VSAN policy requires the ENTERPRISE_PKG license.

You can configure a role so that it only allows commands to be performed for a selected set of VSANs. By default, the VSAN policy for any role is **permit**. In other words, the role can perform commands configured by the rule in all VSANs. In order to selectively allow VSANs for a role, the VSAN policy needs to be set to **deny** and then the appropriate VSANs need to be permitted.

Users configured in roles where the VSAN policy set to **deny** cannot modify configuration for E ports. They can only modify 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, 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. These users cannot perform the following functions that require the startup configuration to be viewed or modified:

- **copy running startup**
- **show startup**
- **show running-config diff**
- **copy startup** <destination>
- **copy** <source> **startup** commands.

For information on these commands, refer to the *Cisco MDS 9000 Family Command Reference*.

Recovering Administrator Password

An administrator can recover a password from a local console connection.

The password recovery procedure must be performed on the supervisor module that becomes the active supervisor module after the recovery procedure is completed. To ensure the other supervisor module does not become the active module, you have two options:

Password recovery is not possible from a Telnet or SSH session.

To recover a administrator password, refer to the *Cisco MDS 9000 Family Command Reference*.

Configuring SSH Services

The Telnet service is enabled by default on all Cisco MDS 9000 Family switches. Before enabling the SSH service, generate a host key pair.

Enabling SSH Service

By default, the SSH service is disabled.

Generating an SSH Host Key Pair

Be sure to have an SSH host key pair with the appropriate version before enabling the SSH service. The SSH service accepts three types of key pairs for use by SSH versions 1 and 2. Generate the SSH host key pair according to the SSH client version used. The number of bits specified for each key pair ranges from 768 to 2048.

- 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.

Using the force Option

If the SSH key pair option is already generated for the required version, use the **force** option to overwrite the previously generated key pair.

About SNMP Security

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 18-2.)

Figure 18-2 SNMP Security

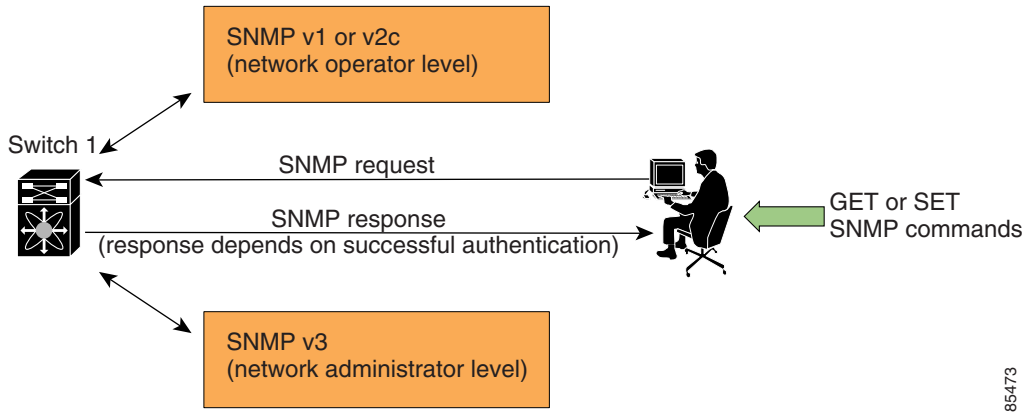

Users and roles configured through the CLI are different from users and roles configured through SNMP. These configurations do not directly correspond with each other. However, you can configure both CLI and SNMP identically, if required. See the "Creating Common Roles" section on page 18-18 for additional information.

SNMP users are different from CLI users. SNMP users also have role-based authentication for roles and authorization purposes.

SNMP Version 1 and Version 2c

SNMPv1 and 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.

Adding a Community String

To add a community string, follows these steps:

Step 1 From Fabric Manager, choose **Security > SNMP** from the Physical pane menu and click the **Communities** tab in the Information pane.

From Device Manager, choose **SNMP > Security** and click the **Communities** tab.

Step 2 Click **Create** on the Device Manager dialog box, or click **Create Row** on the Fabric Manager toolbar.

The Create Community string dialog box displays. (The dialog box from Fabric Manager also provides check boxes to specify one or more switches.)

Step 3 Enter the community name in the Community field.

Step 4 Choose the role from the drop-down list. In Fabric Manager, you can enter a new role name in the field if you do not want to choose one from the list. If you enter a new role name, you must go back and configure this role appropriately (see the "Configuring Common Roles" section on page 18-16).

Step 5 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, follows these steps:

Step 1 From Fabric Manager, select **Security > SNMP** from the Physical pane menu and click the **Communities** tab in the Information pane.

From Device Manager, choose **SNMP > Security** and click the **Communities** tab.

Step 2 Click once to highlight the name of the community you want to delete.

Step 3 Click **Delete** (Device Manager) or the **Delete Row** icon (Fabric Manager).

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.

Adding SNMP Users

To add SNMP users, follows these steps:

Step 1 From Fabric Manager, select **Security > SNMP** from the Physical pane menu and click the **Users** tab in the Information pane.

From Device Manager, choose **SNMP > Security** and click the **Users** tab.

Step 2 Click **Create** on the Device Manager dialog box, or click **Create Row** on the Fabric Manager toolbar.

The Create Users dialog box displays. (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 drop-down list. In Fabric Manager, you can enter a new role name in the field if you do not want to select one from the list. If you enter a new role name, you must go back and configure this role appropriately (see the "Configuring Common Roles" section on page 18-16).

Step 5 Enter the password for the user twice in the New Password and Confirm Password fields.

Step 6 To enable encryption of management traffic, click the **Privacy** check box and complete the password fields.

Enter the authentication password in the Clone Password field to use the same password. Enter a new password twice in the New Password and Confirm Password fields.

Step 7 Click **Create** to create the new entry, or click **Close** to create the entry and close the dialog box.

Deleting SNMP Users

To delete SNMP users, follows these steps:

Step 1 From Fabric Manager, select **Security > SNMP** from the Physical pane menu and click the **Users** tab in the Information pane.

From Device Manager, choose **SNMP > Security** and click the **Users** tab.

Step 2 Click once to highlight the name of the user you want to delete.

Step 3 Click **Delete** (Device Manager) or the **Delete Row** icon (Fabric Manager).

Configuring and Creating SNMP User Roles

To configure users roles, choose **Security > SNMP** from Device Manager, and click the **Roles** tab.

To create a new role, follow these steps:

Step 1 Click **Create**. You see the Create Roles dialog box.

Step 2 Enter an identifier for the role in the Role field.

Step 3 Select one of the following security levels:

- authNoPrv—Authentication without encryption
- AuthPriv—Authentication with encryption

Step 4 For Read access, click the **All** radio button to enable full read access or click **List** and check each check box in the list to enable read access to specific information.

Step 5 For Write access, click the **All** radio button to enable full read access or click **List** and check each check box in the list to enable read access to specific information.

Step 6 Click **Apply** to create the new role, or click **OK** to create the role and close the window.

Viewing SNMP Community and User Information

To view information about SNMP users, roles, and communities from Fabric Manager, choose **Security > SNMP** from the Physical pane menu tree and click the **Users**, **Roles**, or **Communities** tab. You see the list of SNMP users, roles, or communities in the Information pane.

To view this information from the Device Manager, choose **Security > SNMP**. The SNMP dialog box displays.

Group-Based SNMP Access

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 the your user name is created, your roles are set up by your administrator, and you are added to the roles.

Users configured through the CLI are different from users configured through SNMP. These configurations do not directly correspond with each other. However, you can configure both CLI and SNMP identically, if required.

Configuring Common Roles

From Release 1.2(x), 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.

Figure 18-3 Common Roles

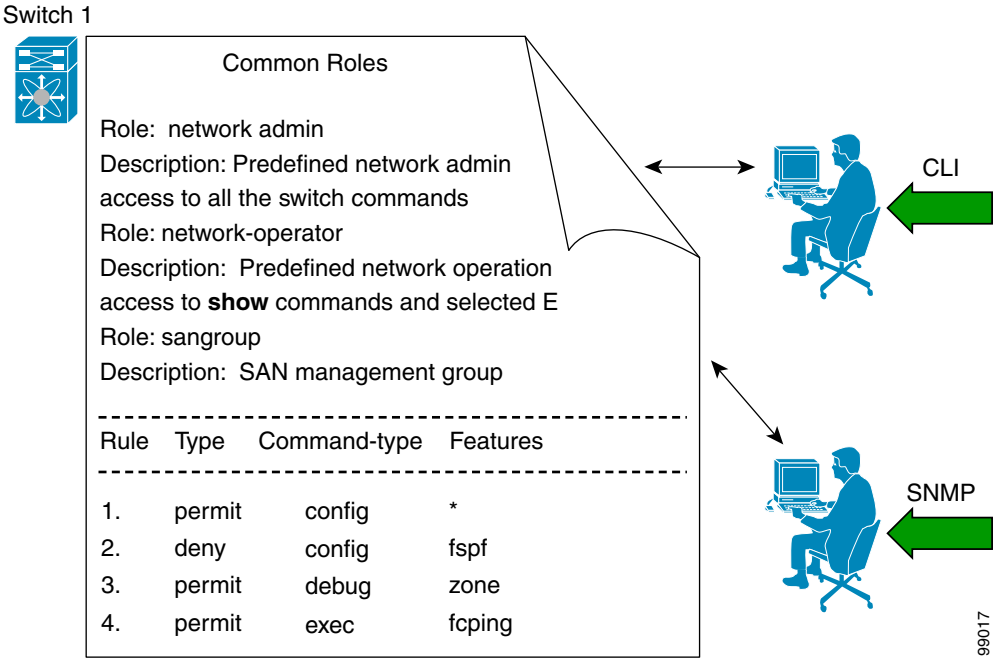

Each role in SNMP is the same as a role created or modified through the CLI. Common Roles allow you to use a set of rules to set the scope of VSAN security. Each role can be restricted to one or more VSANs as required.

You can create new roles or modify existing roles using SNMP or the CLI.

To configure Common Roles from the Device Manager, choose **Common Roles** from the Security menu. You can then access the Rules dialog box to configure the set of rules. To configure Common Roles from Fabric Manager, choose **Security > SNMP** and click the **Roles** tab in the Information pane. Fabric Manager uses a default rules set for roles; therefore, no Rules dialog box is displayed.

See the "Creating Common Roles" section on page 18-18 for additional information.

Creating and Modifying Users

You can create users or modify existing users using SNMP or the CLI.

- SNMP—Create a user as a clone of an existing user in the vsmUserTable on the switch. Once you have created the user, change the cloned secret key before activating the user. Refer to RFC2574.

  You must explicitly configure password(s) for SNMP users. The SNMP user passwords are not generated as the part of the configuration file as they are not portable across devices. The password is limited to a minimum of 8 characters and a maximum of 64 characters.

  An SNMP user must be created on each switch to which the user requires access. If the user is managing 10 switches, each of the 10 switches must have the SNMP user defined.

- CLI—You can 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.

Creating Common Roles

To create a common role in Fabric Manager, perform the following steps:

Step 1 Choose **Security > SNMP** from the Physical pane menu tree, and click the **Roles** tab in the Information pane.

Step 2 Click the **Create Row** icon in the toolbar.

The Roles -> Create dialog box displays.

Step 3 Choose the switches for which you want to configure the 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 Check (or uncheck) the **Has Config and Exec Permission** check box.

If you check the check box, your role will have read, write, and create permission. If you do not check the check box, your role will have read-only permission.

Step 7 Click **Enable** to enable the VSAN scope.

Step 8 Enter the scope in the Scope field.

Step 9 Click **Create** to create the Role, or click **Close** to close the Role dialog without creating the common role.

To create a common role in Device Manager, perform the following steps:

Step 1 Choose **Security > Common Roles**. The Common Roles dialog box displays.

Step 2 Click **Create**.

The Create Common Roles dialog box displays.

Step 3 Enter the name of the common role in the Name field.

Step 4 Enter the description of the common role in the Description field.

Step 5 Click **Enable** to enable the VSAN scope.

Step 6 Enter the scope in the Scope field.

Step 7 Click **Rules** to view the rules for the role, and select the rules you want to enable. Then click **Close** to close the Rules dialog.

The Rules dialog may take a few minutes to display.

Step 8 Click **Create** to create the common role, or click **Close** to close the Common Role dialog without creating the common role.

Editing Common Role Rules (Device Manager Only)

To edit the rules for a common role, perform the following steps:

Step 1 From the Device Manager, choose **Security > Common Roles**.

The Common Roles dialog box displays.

Step 2 Click once on the common role for which you want to edit the rules.

Step 3 Click **Rules** to view the rules for the role.

The Rules dialog may take a few minutes to display.

Step 4 Edit the rules you want to enable or disable for the common role.

Step 5 Click **Apply** to apply the new rules and close the Rules dialog, or click **Close** to close the Rules dialog without applying the rules.

Step 6 Click **Apply** to create the common role, or click **Close** to close the Common Role dialog without creating the common role.

Deleting Common Roles

To delete a common role, perform the following steps:

Step 1 From the Device Manager, choose **Security > Common Roles**. The Common Roles dialog box displays.

From Fabric Manager, choose **Security > SNMP** from the Physical pane menu tree, and click the **Roles** tab in the Information pane.

Step 2 Click once to select the common role you want to delete.

Step 3 Click **Delete** to delete the common role.

Assigning Users to Roles

Once the user and the role are created, the administrator should configure an entry in the vacmSecurityToGroupTable to add the configured user to a configured role.

- To assign users to roles through SNMP, refer to RFC2575.
- To assign users to roles through the CLI, refer to the procedure specified in the *Cisco MDS 9000 Family Command Reference*.

Default Security Settings

Table 18-1 lists the default settings for all security features in any switch.

Table 18-1 Default Security Settings

Parameters	Default
Roles in each switch (for CLI and SNMP users)	Two default roles—network-operator and network-admin.
AAA authentication login	Local authentication is enabled. If the Telnet or SSH options are not specified, the command applies to both.
Telnet server	Enabled.
Accounting log file size on local disk	15,000 bytes.
User's account expiration	Does not expire unless you explicitly configure it to expire.
User name	admin.
User password	admin.
Configured RADIUS sever	Allows access to all RADIUS severs.
RADIUS server timeout interval	The default time-out is one (1) seconds.
RADIUS preshared key	No key is configured.
RADIUS key encryption	clear text (0)—Not encrypted.
RADIUS server connection attempts	A switch tries to connect to a RADIUS server once (1).
RADIUS Authentication port	UDP port 1812.
RADIUS Accounting port	UDP port 1813.
Server key encryption	clear text (0)—Not encrypted.
TACACS+	Disabled
Configured TACACS+ sever	Allows access to all TACACS+ severs.
TACACS+ server timeout interval	The default time-out is one (5) seconds.
TACACS+ preshared key	No key is configured.
TACACS+ key encryption	clear text (0)—Not encrypted
TACACS+ server connection attempts	A switch tries to connect to a TACACS+ server once (1).
TACACS+ Authentication port	UDP port 49.
VSAN policy	Permit.

Restricting Switch Access

You can restrict access to a Cisco MDS 9000 Family switch using IP Access Control Lists (IP-ACLs).

Configuring Fabric Security

Fibre Channel Security Protocol (FC-SP) capabilities in SAN-OS 1.3(x) provides 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 implemented in SAN-OS1.3(x) to provide authentication between Cisco MDS switches and other devices. It consists of the CHAP protocol combined with the Diffie-Hellman exchange.

This chapter contains the following topics:

- About Fabric Authentication, page 19-1
- About DHCHAP, page 19-2
- DHCHAP Compatibility with Existing MDS Features, page 19-2
- Configuring DHCHAP Authentication, page 19-3
- Enabling DHCHAP, page 19-3
- Configuring DHCHAP Authentication Modes, page 19-3
- Configuring the DHCHAP Hash Algorithm, page 19-4
- Configuring DHCHAP Groups, page 19-4
- Configuring DHCHAP Passwords, page 19-4
- Configuring Passwords for Other Devices, page 19-5
- Configuring the DHCHAP Timeout Value, page 19-5
- Default Fabric Security Settings, page 19-5

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 switches and hosts 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.

Figure 19-1 Authentication between Switches and Hosts

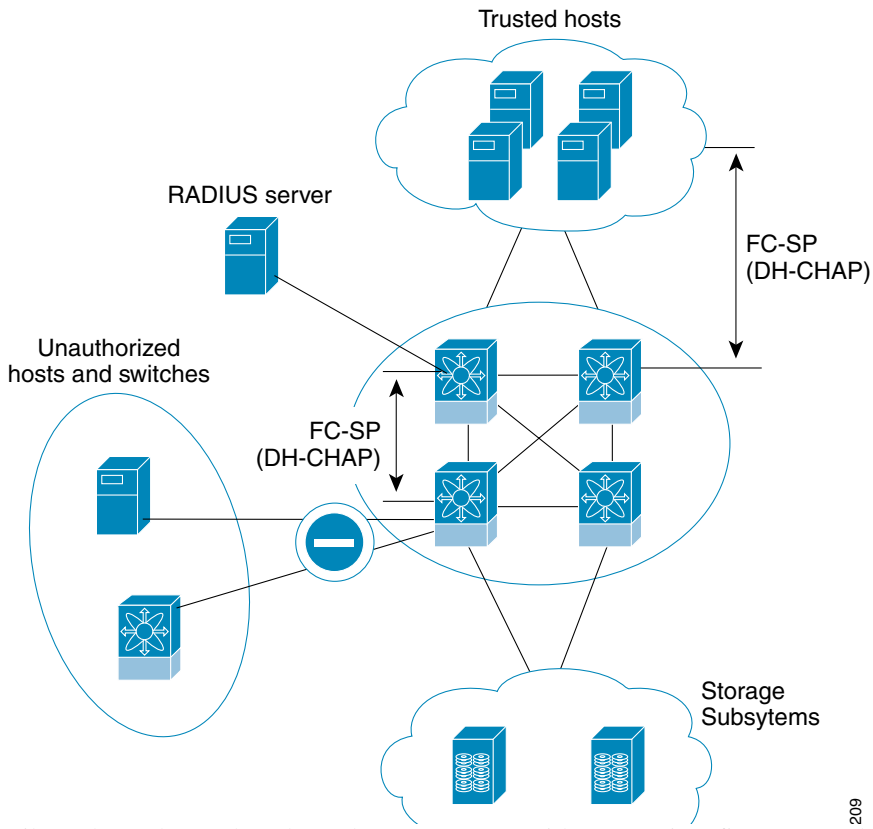

Fibre Channel (FC) host bus adapters (HBAs) with appropriate firmware and drivers are required for host-switch authentication.

About 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.

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 MD-5 and SHA-1 algorithm-based authentication.

Configuring the DHCHAP feature requires the ENTERPRISE_PKG license.

DHCHAP Compatibility with Existing MDS Features

This sections identifies the impact of configuring the DHCHAP feature along with existing 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 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 per-VSAN basis.
- High availability--DHCHAP authentication works transparently with existing HA features.

Configuring DHCHAP Authentication

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 password for the local switch and other switches in the fabric.

Step 5 Configure the timeout value for reauthentication.

Step 6 Verify the DHCHAP configuration.

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.

Configuring 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—Does not support DHCHAP authentication. Authentication messages sent to such ports return error messages to the initiating switch.

Whenever DHCHAP port mode is changed to a mode other than the Off mode, reauthentication is performed.

Table 19-1 identifies the switch-to-switch authentication behavior between two Cisco MDS switches in various modes.

Table 19-1 DHCHAP Authentication Status Between Two MDS Switches

Switch N	**Switch 1**
DHCHAP Modes	**on**
on	FC-SP authentication is performed
auto-Active	FC-SP authentication is not performed.
auto-Passive	FC-SP authentication is not performed.
off	Link is brought down

Configuring the DHCHAP Hash Algorithm

Cisco MDS switches support a default hash algorithm priority list of MD-5 followed by SHA-1 for DHCHAP authentication.

If you change the hash algorithm configuration, ensure to change it globally for all switches in the fabric.

RADIUS and TACACS+ protocols always use MD-5 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 DHCHAP Groups

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.

If you change the DH group configuration, ensure to change it globally for all switches in the fabric.

Configuring DHCHAP Passwords

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 which participate in DHCHAP:

- Approach 1—Use the same password for all switches in the fabric--the simplest approach. When you add a new switch, you will use the same password to authenticate that switch in this fabric. It is also the most vulnerable approach if someone from 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.

We recommend using RADIUS or TACACS+ for fabrics with more than five switches. If you need to use 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. Refer to the Cisco MDS 9000 Family Fabric Manager User Guide for further information.

All passwords are restricted to 64 alphanumeric characters and can be changed, but not deleted.

Configuring Passwords for Other 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 know 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).

The switch WWN identifies the physical switch. This WWN is used to authenticate the switch and is different from the VSAN node WWN.

Configuring the DHCHAP Timeout 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 timeout value consider the following factors:

- The existing RADIUS and TACACS+ timeout values.
- The same value must also be configured all switches in the fabric.

Default Fabric Security Settings

Table 19-2 lists the default settings for all fabric security features in any switch.

Table 19-2 Default Fabric Security Settings

Parameters	Default
DHCHAP feature	Disabled.
DHCHAP hash algorithm	A priority list of MD-5 followed by SHA-1 for DHCHAP authentication
DHCHAP authentication mode	auto-passive.

Table 19-2 Default Fabric Security Settings (continued)

DHCHAP group default priority exchange order	0, 4, 1, 2, and 3 respectively.
DHCHAP timeout value	30 seconds.

Configuring Port Security

All switches in the Cisco MDS 9000 Family provide port security features that rejects intrusion attempts and reports these intrusions to the administrator.

Note Port security is only supported for Fibre Channel ports.

This chapter contains the following topics:

- Port Security Features, page 20-1
- About Auto-Learn, page 20-3
- Manually Configuring Port Security, page 20-7
- Database Scenarios, page 20-8
- Displaying Port Security Statistics, page 20-9
- Displaying Port Security Violations, page 20-9
- Default Port Security Settings, page 20-9

Port Security Features

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 syslog messages.

Enforcing Port Security

To enforce port security, configure the devices and switch port interfaces through which each device or switch is connected.

- 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 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.

Configuring a Port Binding

To configure a port binding on a switch, perform the following steps.

Step 1 From the Fabric Manager, choose **Port Security** from one of the VSANs on the menu tree.

The Information pane of the Fabric Manager displays port security information for that VSAN.

Step 2 Click the **Config** tab.

You see a list of the port security configured port bindings for that VSAN.

Step 3 Click the **Create Row** icon.

The Create Binding dialog box displays.

Step 4 Choose the switch for which you want to create the port binding from drop-down list.

Step 5 Choose the WWN DEVICE device type for that switch.

Step 6 Enter the PORT ID of the switch to bind to.

Step 7 Enter the port type.

Step 8 Enter the interface (e.g. fc1/1)

Step 9 Click **Create** to creating the port binding, or click **Close** to close the Create Binding dialog box without creating a port binding.

Copying an Active Configuration to the Running Configuration

To copy the active configuration to the running configuration, perform the following steps.

Step 1 From the Fabric Manager, choose **Port Security** from one of the VSANs on the menu tree.

The Information pane of the Fabric Manager displays Port Security information for that VSAN.

Step 2 Click the **Action** tab.

You see a list of switches for that VSAN.

Step 3 Check the **CopyActive ToConfig** check box next to the switch for which you want to copy the configuration.

The active configuration is copied to the running configuration when the binding is activated.

Step 4 Uncheck the check box if you do not want the configuration copied when the binding is activated.

Deleting a Port Binding

To delete a port binding on a switch, perform the following steps.

Step 1 From the Fabric Manager, choose **Port Security** from one of the VSANs on the menu tree.

The Information pane of the Fabric Manager displays Port Security information for that VSAN.

Step 2 Click the **Config** tab.

You see a list of the port security configured port bindings for that VSAN.

Step 3 Click the row you want to delete.

Step 4 Click the **Delete Row** icon.

You see a confirmation dialog box.

Step 5 Click **Yes** to delete the row, or click **No** to close the dialog box without deleting the row.

About Auto-Learn

You can instruct the switch to automatically learn (auto-learn) the port security configurations over a specified period. The **auto-learn** option 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 port security feature for the first time as it saves tedious manual configuration for each port. Auto-learn is configured 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 a **shutdown** command is issued on that port.

Activating Port Security

By default, the port security feature is not activated.

When you activate the port security feature, the **auto-learn** option is also automatically enabled. You can choose to activate the port-security feature and disable auto-learn. In this case, you need to manually populate the port security database by individually securing each port.

Activating a Port Binding

To activate a port security port binding, perform the following steps.

Step 1 From the Fabric Manager, choose **Port Security** from one of the VSANs on the menu tree.

The Information pane of the Fabric Manager displays Port Security information for that VSAN.

Step 2 Click the **Action** tab.

You see a list of switches for that VSAN.

Step 3 Click in the Action column under Activation, next to the switch for which you want to activate a port binding.

Step 4 Choose the port binding option that you want to apply to the switch from the drop-down menu. Choose from the following options:

- Activate—Valid port bindings are activated.
- Activate (TurnLearningOff)—Valid port bindings are activated and the **auto-learn** option is turned off.
- ForceActivate—Activation is forced.
- ForceActivate (TurnLearningOff)—Activation is forced and the **auto-learn** option is turned off.
- Deactivate—Deactivates all currently active port bindings.
- NoSelection—No action is taken

Displaying Activated Port Bindings

To display port security active port bindings, perform the following steps.

Step 1 From the Fabric Manager, choose **Port Security** from one of the VSANs on the menu tree.

The Information pane of the Fabric Manager displays port security information for that VSAN.

Step 2 Click the **Active** tab.

You see a list of the port security active port bindings for that VSAN.

Configuring 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, the **auto-learn** option is disabled by default.
- If the port security feature is activated, the **auto-learn** option is enabled by default.

If the **auto-learn** option is enabled on a VSAN, you cannot activate the database for that VSAN without the **force** option.

Table 20-1 summarizes the authorized connection for device requests.

Table 20-1 Auto-learn Device Authorization

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	Aport that is not configured	Permitted if **auto-learn** option enabled	3
Denied if **auto-learn** disabled	4		

Table 20-1 Auto-learn Device Authorization (continued)

Configured or not configured	A switch port that allows any device	Permitted	5
Configured to login to any switch port	Any port on the switch	Permitted	6
Not configured	A port configured with some other device	Denied	7

Authorization Scenario

Assuming 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 20-2 summarizes the port security authorization results for this active database.

Table 20-2 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
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

Table 20-2 Authorization Results for Scenario (continued)

Scenario	Device Connection Request	Authorization	Condition	Reason
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

Turning Auto-Learning On or Off

To turn Auto-learning on or off, perform the following steps.

Step 1 From the Fabric Manager, choose **Port Security** from one of the VSANs on the menu tree.

The Information pane of the Fabric Manager displays Port Security information for that VSAN.

Step 2 Click the **Action** tab.

You see a list of switches for that VSAN.

Step 3 Click in the AutoLearn column next to the switch for which you want to enable AutoLearning.

Step 4 Choose **on** from the drop-down menu to turn on AutoLearning; choose **off** to turn off AutoLearning for that switch.

Manually Configuring Port Security

To configure port security in 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.

Identifying WWNs to Configure Port Security

If you decide to manually configure port security, be sure to adhere to the following guidelines:

- Identify switch ports by the interface or the fWWN.
- Identify devices by the pWWN or nWWN.
- If an Nx port:
  - is allowed to login to SAN switch port Fx, then that Nx port can only log in through the specified Fx port.
  - nWWN is bound to a 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.
- Saving the running configuration saves the configuration database and activated entries in the active database. Learned entries in the active database are not saved.

Securing Authorized Ports

After identifying the WWN pairs that need to be bound, add those pairs to the port security database.

Activating the Port Security Database

When you activate the port security 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. Additionally, all devices that have already logged into the VSAN at the time of activation are also learned and added to the active database. If the **auto-learn** option is already enabled in a VSAN, you will not be allowed to activate the database.

Database activation is rejected in the following cases:

- Missing or conflicting entries exist in the configuration database but not in the active database.
- The **auto-learn** option was enabled before the activation.
- The exact security is not configured for each PortChannel member.
- If the configured database is empty and the active database is not.

Forcing Port Security 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.

An activation using the **force** option logs out existing devices if they violate the active database.

Reactivating the Database

If the **auto-learn** option is enabled and you activate the database, you will not be allowed to proceed.

Database Scenarios

Table 20-3 lists the differences and interaction between the active and configuration databases.

Table 20-3 Active and Configuration Port Security Databases

Configuration Database	Active Database
Read-write.	Read only.
Saving the configuration saves all the entries in the configuration database.	Saving the configuration only saves the activated entries. Learned entries are not saved.
Once activated, the configuration database can be modified without any effect on the active database.	Once activated, all devices that have already logged into the VSAN are also learned and added to the active database.
You can overwrite the configuration database with the active database.	You can overwrite the active database with the configured database by activating the port security database. An activation using the **force** option may violate the entries already configured in the active database.

Displaying Port Security Statistics

To display port security statistics, perform the following steps.

Step 1 From the Fabric Manager, choose **Port Security** from one of the VSANs on the menu tree.

The Information pane of the Fabric Manager displays port security information for that VSAN.

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, perform the following steps.

Step 1 From the Fabric Manager, choose **Port Security** from one of the VSANs on the menu tree.

The Information pane of the Fabric Manager displays port security information for that VSAN.

Step 2 Click the **Violations** tab.

You see a list of the port security violations for that VSAN.

Default Port Security Settings

Table 20-4 lists the default settings for all security features in any switch.

Table 20-4 Default Security Settings

Parameters	Default
Auto-learn	Enabled if port security is enabled
Port security	Disabled

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.
  - FSPF 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.

This chapter contains the following topics:

- FSPF Features, page 21-2
- FSPF Examples, page 21-2
- Configuring FSPF Globally, page 21-4
- Configuring FSPF for a Specific Interface, page 21-5
- Configuring Fibre Channel Routes, page 21-7
- Broadcast Routing, page 21-8
- In-Order Delivery, page 21-8
- Configuring Flow Statistics, page 21-10
- Viewing FSPF Statistics, page 21-10
- Default Settings, page 21-10

FSPF Features

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 re convergence 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. The FSPF feature can be used on any topology.

Fault Tolerant Fabric

Figure 21-1 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 21-1 Fault Tolerant Fabric

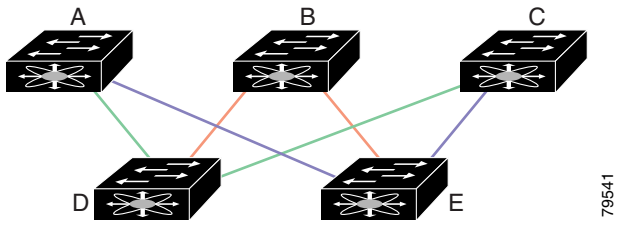

For example, if all links are of equal speed, the FSPF calculates two equal paths from A to C: A-D-C (green) and A-E-C (blue).

Redundant Links

To further improve on the topology in, each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. 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 21-2 Fault Tolerant Fabric with Redundant Links

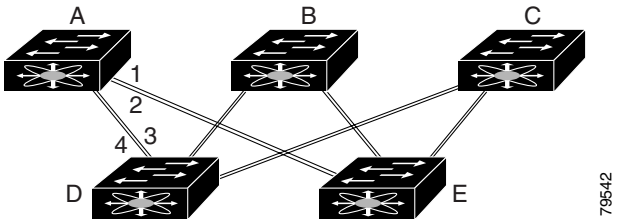

For example, if all links are of equal speed and no PortChannels exist, the FSPF calculates four equal paths from A to C: A1-E-C, A2-E-C, A3-D-C, and A4-D-C. If PortChannels exist, these paths are reduced to two.

Fail-over Scenarios for PortChannels and FSPF Links

The SmartBits traffic generator was used evaluate the scenarios displayed in Figure 21-3 and summarized in Table 21-1 and Table 21-2. 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 1G in two scenarios:

- Disabling the traffic link by either physically removing the cable.
- Shutting down either switch 1 or switch 2.

Figure 21-3 Fail-over Scenario Using Traffic Generators

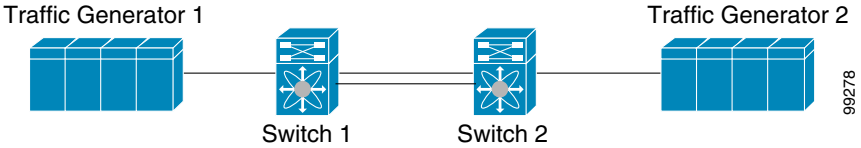

Table 21-1 Physically Removing the Cable for the SmartBits Scenario

PortChannel Scenario	FSPF Scenario (Equal cost ISL)
Switch 1	**Switch 2**
110ms (~2K frame drops)	130ms+ (~4k frame drops)
100 ms hold time when a signal loss is reported as mandated by the standard	

Table 21-2 Shutting Down the Switch for the SmartBits Scenario

PortChannel Scenario	FSPF Scenario (Equal cost ISL)
Switch 1	**Switch 2**
~0ms (~8 frame drops)	110ms (~2K frame drops)
No hold time needed.	Signal loss on switch 1

Configuring FSPF Globally

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.

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.

Managing FSPF General Attributes

To manage FSPF general attributes, perform the following steps:

Step 1 From the Fabric Manager, choose **FC > FSPF** on the menu tree and click the **General** tab.

From the Device Manager, choose **FC > FSPF** and click the **General** tab.

The Information pane from the Fabric Manager displays information for multiple switches. The dialog box from the Device Manager displays FSPF information for a single switch.

Step 2 Configure the FSPF general attributes.

Disabling FSPF Routing Protocols

By default, FSPF is enabled on switches in the Cisco MDS 9000 Family.

Link State Record Defaults

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. displays the default settings for switch responses.

Table 21-3 LSR Default Settings

LSR Option	Default	Description
Acknowledgement interval (RxmtInterval)	5 seconds	The time a switch waits for an acknowledgement 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.

Viewing Link State Records

To monitor FSPF LSRs from the Device Manager, choose **FC > FSPF** and click the **LSDB LSRs** tab.

Viewing FSPF Links

To view FSPF links from the Device Manager, choose **FC > FSPF** and click the **LSDB Links** tab.

Configuring FSPF for a Specific Interface

Several FSPF commands are available on a per interface basis. The following configuration procedures apply to an interface in a specific VSAN and are described in this section.

This section contains the following topics:

- Configuring FSPF Interfaces, page 21-6
- Computing Route Cost, page 21-6
- Specifying Hello Time Intervals, page 21-6
- Specifying Dead Intervals, page 21-6
- Disabling FSPF for Specific Interfaces, page 21-6
- Retransmitting Intervals, page 21-6
- Viewing FSPF Interface Statistics, page 21-7

Configuring FSPF Interfaces

To configure FSPF interfaces, perform the following steps:

Step 1 From the Fabric Manager, choose **FC > FSP**F on the menu tree and click the **Interfaces** tab.

To configure FSPF interfaces from the Device Manager, choose **FC > FSPF** and click the **Interfaces** tab.

Step 2 Configure the attributes for the FSPF interfaces.

Computing Route 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 seconds. The default cost for 1Gbps is 1000 and 2Gbps is 500 seconds

Specifying 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.

This value must be the same in the ports at both ends of the ISL.

Specifying Dead 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.

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.

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.

FSPF must be enabled at both ends of the interface for the protocol to work.

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.

This value must be the same on the switches on both ends of the interface.

Viewing FSPF Interface Statistics

To monitor FSPF interface statistics from the Fabric Manager, choose **FC > FSPF** on the menu tree and click the **Interface Stats** tab.

To monitor FSPF interface statistics from the Device Manager, choose **FC > FSPF** and click the **Interface Stats** tab.

Configuring 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.

Figure 21-4 Fibre Channel Routes

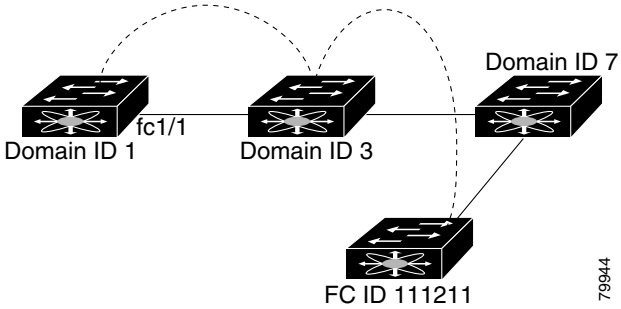

Other than in VSANs, run time checks are not performed on configured and suspended static routes.

To configure Fibre Channel routes, perform the following steps:

Step 1 From the Device Manager, choose **FC > Routes**. The dialog box displays routes for a single switch.

Step 2 Configure the attributes for the route.

Step 3 To add a route from Device Manager, click **Create** in the dialog box.

You see the Create Route dialog box.

Step 4 Click the button to the right of the Interface field and select the interface on which to configure the Fibre Channel route.

Step 5 Complete the other fields on this window and click **OK** to add a route.

Configuring Fibre Channel Route Flows

To view Fibre Channel flows and add a route flow, perform the following steps:

Step 1 From the Fabric Manager, choose **FC > Route Flow Statistics** on the menu tree. The Information pane from Fabric Manager displays flows for multiple switches.

From the Device View, choose **FC > Routes** and click the **Flow Statistics** tab. The dialog box from the Device Manager displays flows for a single switch.

Step 2 Configure the flow attributes for the route.

Step 3 To add a route flow from Fabric Manager, click **Create Row** on the toolbar.

To add a route flow from Device Manager, click **Create** in the dialog box.

You see the Create Route Flow dialog box.

Step 4 Complete the fields in this dialog box and click **Create** to add a route flow.

Broadcast Routing

Broadcast in a Fibre Channel fabric uses the concept of a distribution tree to reach all switches in the fabric (for broadcast traffic).

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 the distribution tree information. The protocols create a loop-free broadcast distribution tree.

Caution All switches in the fabric should run the same multicast and broadcast distribution tree algorithm to ensure the same distribution tree.

In-Order Delivery

In-order delivery 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.

In case of a single switch, 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.

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.

Figure 21-5 Route Change Delivery

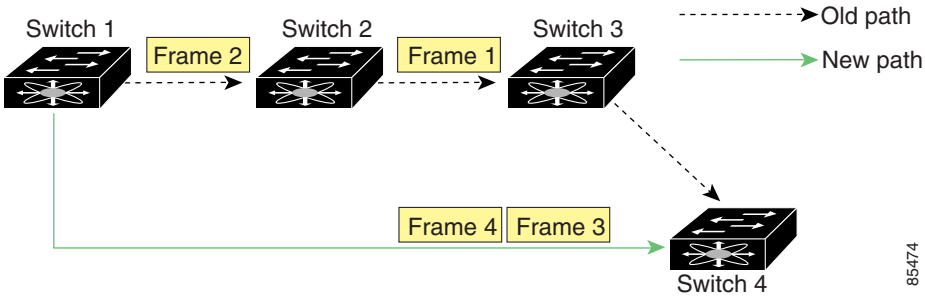

In Figure 21-5, the new path from Switch 1 to Switch 4 is faster. Hence, Frame 3 and Frame 4 may be delivered before Frame 1 and Frame 2.

If the in-order guarantee feature is enabled, the frames within the network are treated as specified below:

- Frames in the network are delivered in the order in which they are transmitted.
- Frames which can not be delivered in-order within the network latency drop period are dropped inside the network.
- The number of dropped frames are reduced by slowing down the traffic at the frame source.

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.

Figure 21-6 Link Congestion Delivery

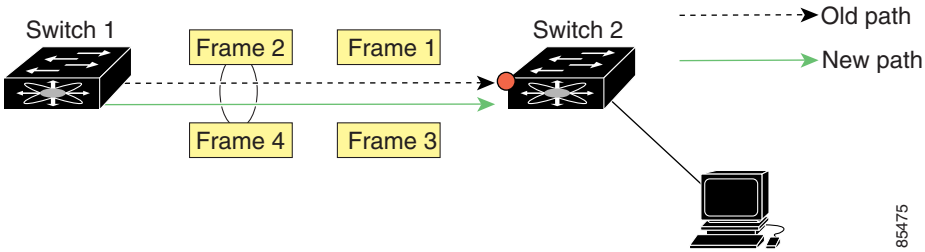

In Figure 21-6, the port of the old path (red dot) is congested. Hence Frame 3 and Frame 4 can be delivered before Frame 1 and Frame 2.

When the in-order guarantee feature is enabled, the frames crossing a PortChannel are treated as specified below:

- Frames using the old path are delivered before new frames are accepted.
- Frames which cannot be delivered in-order, through the old path, within the switch latency drop period are dropped.
- The new frames are delivered through the new path after the switch latency drop period has elapsed.

Enabling In-Order Delivery

By default, in-order delivery is disabled on switches in the Cisco MDS 9000 Family.

Tip We recommend that you only enable this feature in a switch when devices are present in the switch that cannot handle any out-of-order frames. 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 due to an intentional pausing of fabric forwarding to purge the fabric of resident frames that could potentially be forwarded out of order.

Configuring Flow Statistics

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

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.

Viewing FSPF Statistics

To monitor FSPF statistics from the Fabric Manager, choose **FC > FSPF** on the menu tree and click the **Statistics** tab.

To monitor FSPF statistics from the Device Manager, choose **FC > FSPF** and click the **Statistics** tab.

Default Settings

Table 21-4 lists the default settings for FSPF features.

Table 21-4 Default FSPF Settings

Parameters	Default
FSPF	Enabled on all E ports and TE ports.
SPF computation	Dynamic.
SPF hold time	0.
Backbone region	0.
Acknowledgement interval (RxmtInterval)	5 seconds.
Refresh time (LSRefreshTime)	30 minutes.

Table 21-4 Default FSPF Settings (continued)

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.

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 or 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.

This chapter contains the following topics:

- Traffic Management Services, page 22-2
- Configuring the Ethernet Management Port, page 22-2
- Configuring the Default Gateway, page 22-3
- Configuring the Default Network, page 22-4
- IP Access Control Lists, page 22-4
- Configuring IPFC, page 22-8
- Configuring Overlay VSANs, page 22-8
- Configuring Multiple VSANs, page 22-9
- Configuring VRRP, page 22-10
- Default Settings, page 22-14
- Managing IPFC Connectivity with Multiple VSANs, page 22-10
- Enabling or Disabling IP Forwarding, page 22-15
- Viewing Information and Statistics, page 22-15

Traffic Management Services

In-band options are compliant with and use the RFC 2625 standards. An NMS host running IP protocol over a 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.

Figure 22-1 Management Access to Switches

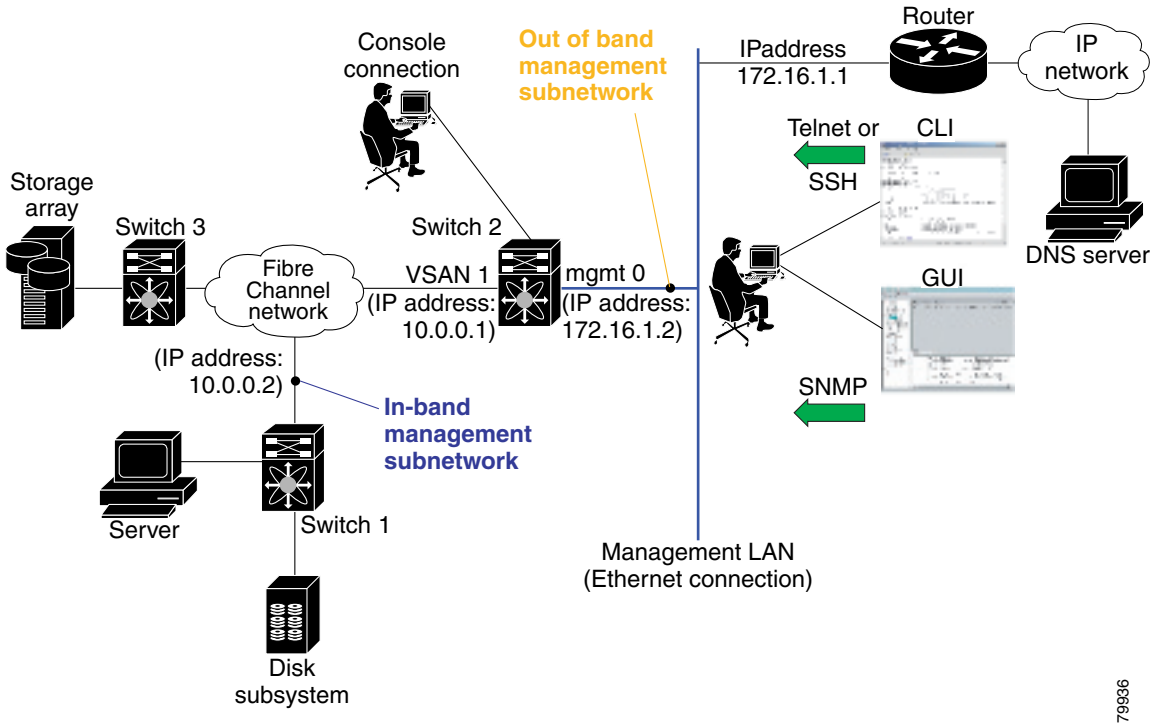

Configuring the Ethernet Management Port

The management port on the switch allows multiple simultaneous Telnet or SNMP network management sessions. You can also configure the supervisor module Ethernet interface and VSAN interfaces as management ports. This section focuses on the Ethernet management port (mgmt0). You can remotely configure the switch through the management port. To configure a connection remotely, you must configure the IP parameters (IP address and subnet mask) from the CLI so that the switch is reachable.

Before you begin to configure the management interface manually, obtain the switch IP address and IP subnet mask. Also make sure the console cable is connected to the console port.

Configuring the Default Gateway

The IP address for a switch's default gateway should be configured along with the IP static routing commands (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 IP addresses will be used regardless of the default-gateway being enabled or disabled. If these IP address are configured and 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.

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 IP address of the gateway switch.

Figure 22-2 Overlay VSAN Functionality

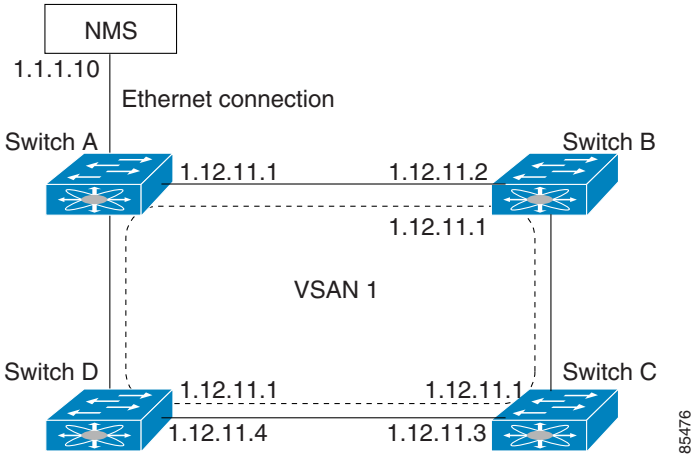

In Figure 22-2, Switch A has the IP address 1.12.11.1, Switch B has the IP address 1.12.11.2, Switch C has the IP address 1.12.11.3, and Switch D has the IP address 1.12.11.4. Switch A is the gateway switch with the Ethernet connection. The NMS uses the IP address 1.1.1.10 to connect to the gateway switch. Frames forwarded to any switch in the overlaid VSAN 1 are routed through the gateway switch. Configuring the gateway switch's IP 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.

Configuring the Default Network

When IP routing is enabled on the switch, assign the IP default network address. The switch considers routes to that network as the last resort. If the IP default network address is not available, the switch uses the IP default gateway address. For every network configured with the IP 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 IP addresses will be used regardless of the default gateway being enabled or disabled. If these IP address are configured and not available, the switch will fall back to using the default gateway IP address, if you have configured it. Be sure to configure IP addresses for all entries in the switch.

Configuring an IP Route

To configure an IP route or identify the default gateway, perform the following steps.

Step 1 From the Device Manager, choose **IP > Routes**.

You see the IP Routes window.

Step 2 To create a new IP route or identify the default gateway on a switch, click **Create**.

You see the Create IP Routes window.

Step 3 Complete the fields in this window and click **OK** to add an IP route.

Step 4 To configure a static route, enter the destination network ID and subnet mask in the Dest and Mask fields. To configure a default gateway, enter the IP address of the seed switch in the Gateway field.

IP Access Control Lists

IP Access control lists (IP-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IP-ACLs restrict IP-related MDS out-of-band management traffic and in-band traffic based on IP addresses (Layer 3 and Layer 4 information).

You can use IP-ACLs to control transmissions on an interface.

IP-ACL Configuration Guidelines

Follow these guidelines when configuring IP-ACLs in any switch or director in the Cisco MDS 9000 Family:

- IP-ACLs cannot be configured for Gigabit Ethernet interfaces (IPS modules) or for Fibre Channel interfaces.
- IP-ACLs can only be configured on the management interface and VSAN interfaces.

- An IP-ACL is a sequential collection of permit and deny conditions that apply to IP flows. Each IP packet is tested against the conditions in the list. The first match determines if the software accepts or rejects the rule. Because the software stops testing conditions after the first match, the order of the conditions in the list is critical. If no conditions match, the software rejects that rule.
- An IP protocol can be configured using an integer ranging from 0 to 255 to represent a particular IP protocol. Alternatively, you can specify the name of a protocol: **icmp, ip, tcp**, or **udp.** IP includes Transmission Control Protocol (TCP), User Datagram Protocol (UDP), Internet Control Message Protocol (ICMP), and other protocols.
- The source/source-wildcard and destination/destination-wildcard is specified 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).
  - Using the **any** option as an abbreviation for a source/source-wildcard or destination/destination-wildcard (0.0.0.0/255.255.255.255)

To configure an IP-ACL, you must complete the following tasks:

- Create an IP-ACL by specifying a name and access condition.

  All lists use the source and destination address for matching operations. You can configure finer granularity using optional keywords
- Apply the access list to specified interfaces.

Creating IP-ACLs

You can specify IP- ACLs using a assigned name. Each IP-ACL can have a maximum of 256 entries. Each entry is a unique filter applied to a specified interface. Each switch can have a maximum of 64 IP-ACLs.

Traffic coming into the switch is compared to IP-ACL entries based on the order that the entries occur in the switch. New statements are added to the end of the list. The switch keeps looking until it has a match. If no matches are found when the switch reaches the end of the list, the traffic is denied. For this reason, you should have the frequently hit entries at the top of the list. There is an implied deny for traffic that is not permitted. A single-entry IP-ACL with only one **deny** entry has the effect of denying all traffic.

Adding Entries to an Existing IP-ACL

After you create an IP- ACL, you place subsequent additions at the end of the IP-ACL. You cannot insert entries in the middle of an IP-ACL. Each configured entry is automatically added to the end of a IP-ACL.

Comparing Ports

Use the following operators to compare the source and destination ports:

- eq = equal
- gt = greater than
- lt = less than
- range = range of ports

Port numbers range from 0 to 65535 for TCP and UDP ports. displays the port numbers for associated TCP and UDP ports.

Table 22-1 TCP and UDP Port Numbers

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	
radius	1812	
wbem-http	5988	
wbem-https	5989	
UDP	dns	53
tftp	69	
ntp	123	
snmp	161	
snmp-trap	162	
syslog	514	

ICMP packets are filtered by the ICMP message type or the message code. Both values range from 0 to 255. displays the value for each associated ICMP type.

Table 22-2 ICMP Type Value

ICMP Type	Value
echo	8
echo-reply	0
destination unreachable	3
traceroute	30
time exceeded	11

Applying IP-ACLs

You can define IP-ACLs without applying them. However, the IP-ACLs will have no effect until they are applied to the switch's interface.

Tip Apply the IP-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 IP-ACL to E0 on Switch 1 instead of an outbound list to E1 on Switch 3.

Figure 22-3 Denying Traffic on the Inbound Interface

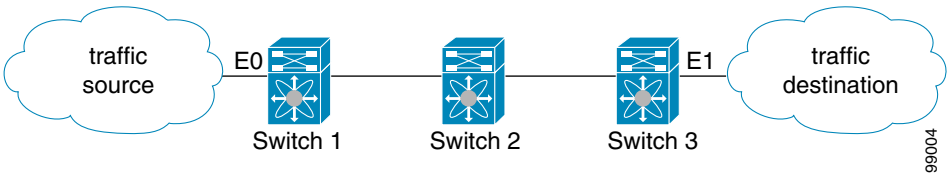

The access group controls access to an interface. Each interface can only be associated with one access list per direction. The ingress direction can have a different ACL than the egress direction. The access group becomes active on creation.

Tip We recommend creating all rules in an access list, before creating the access group that uses this access -list.

Caution If you create an access group before an access-list, all packets in that interface are dropped, because the access list is empty.

The terms in, out, source, and destination are used as referenced by the switch.

- In—Traffic that is arriving on the interface and which will go through the switch; the source would be where it's been and the destination is where it's going (on the other side of the router).

The access-group configuration for the ingress traffic applies to both local and remote traffic.

- Out—Traffic that has already been through the switch and is leaving the interface; the source would be where it's been (on the other side of the router) and the destination is where it's going.

The access-group configuration for the egress traffic applies only to local traffic.

Using the **log-deny** option at the end of the individual ACL entries shows the ACL number and whether the packet was permitted or denied, in addition to port-specific information. This option causes an information logging message about the packet that matches the dropped entry (or entries).

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 dumped to the log.

The following is an example of 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 is an example of 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
```

Configuring IPFC

Once the VSAN interface is created, you can specify the IP address for that VSAN.

Configuring 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.

To configure an overlay VSAN, follow these steps:

To configure the management interface displayed in Figure 22-4, set the default gateway to an IP address on the Ethernet network.

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 default gateway (route) and the IP address on switches that point to the NMS.

Figure 22-4 Overlay VSAN Configuration Example

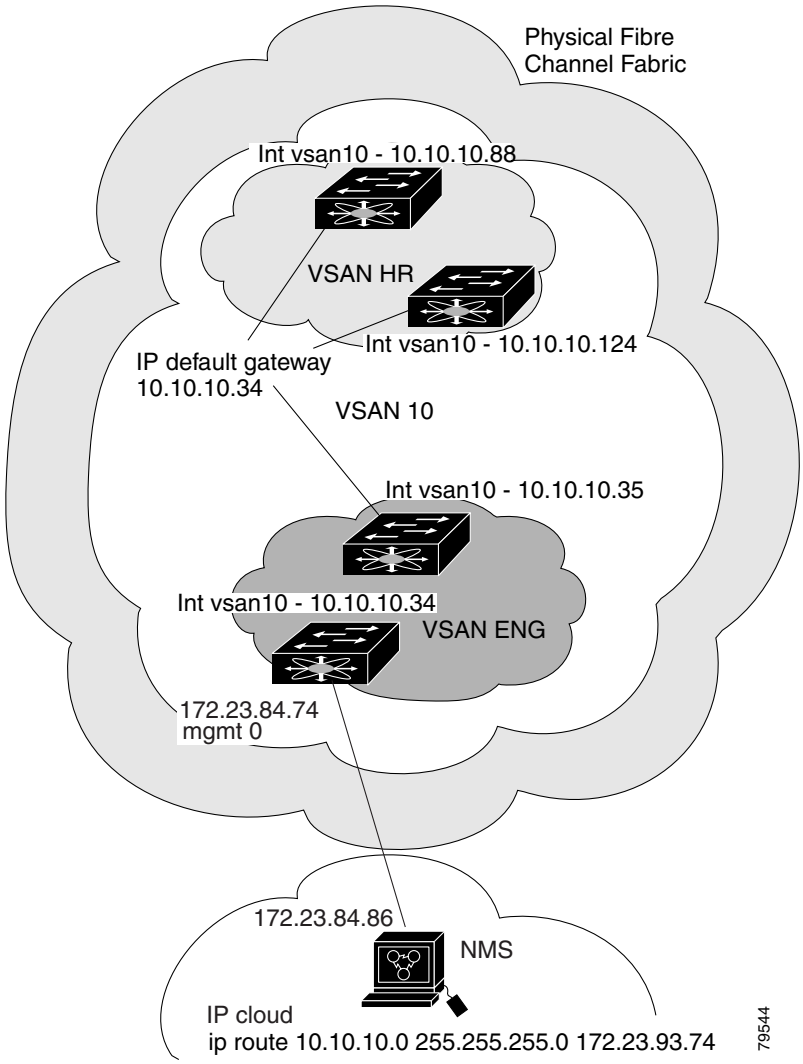

Configuring Multiple VSANs

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

To configure an overlay VSAN, 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 route on the Fibre Channel switches and the IP cloud.

Figure 22-5 Multiple VSANs Configuration Example

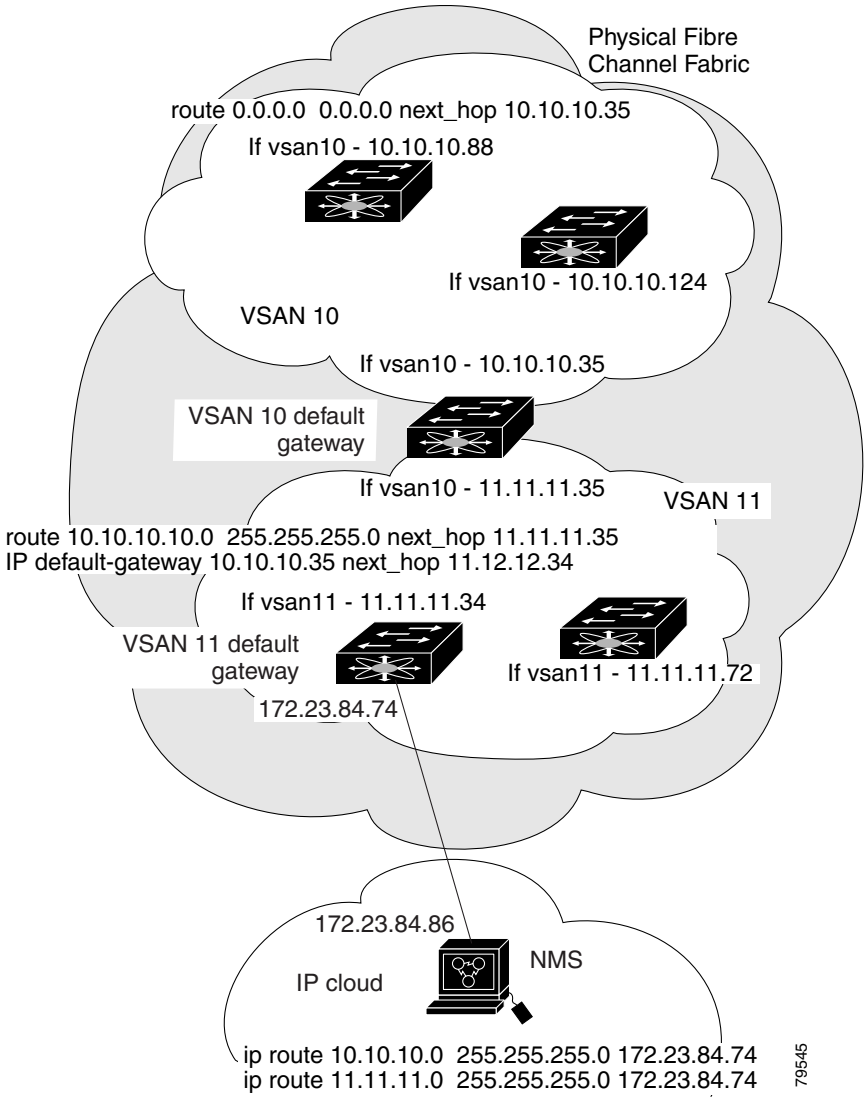

Managing IPFC Connectivity with Multiple VSANs

To configure IPFC from the Device Manager, choose VSAN from the FC menu and click the General tab.

Configuring VRRP

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.

VRRP Features

VRRP provides a redundant alternative path to the gateway switch, which has connectivity to the NMS. VRRP has the following characteristics and advantages:

- VRRP is a restartable application.
- When a VRRP master fails, the VRRP backup takes over within three times the advertisement time.
- VRRP over Ethernet, VRRP over VSAN, and Fibre Channel functions are implemented as defined in RFC 2338.
- A virtual router is mapped to each VSAN and Ethernet interface with its unique virtual router IP, virtual router MAC, and VR ID.
- VR IDs can be reused in multiple VSANs with a 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.

VRRP Functionality

In Figure 22-6, switch A is the VRRP master and switch B is the VRRP backup switch. Both switches have IP address to VRRP mapping configured. The other switches set switch A as the default gateway. If switch A fails, the other switches don't have to change the routing configurations as switch B automatically becomes the master and takes over the function of a gateway.

Figure 22-6 VRRP Functionality

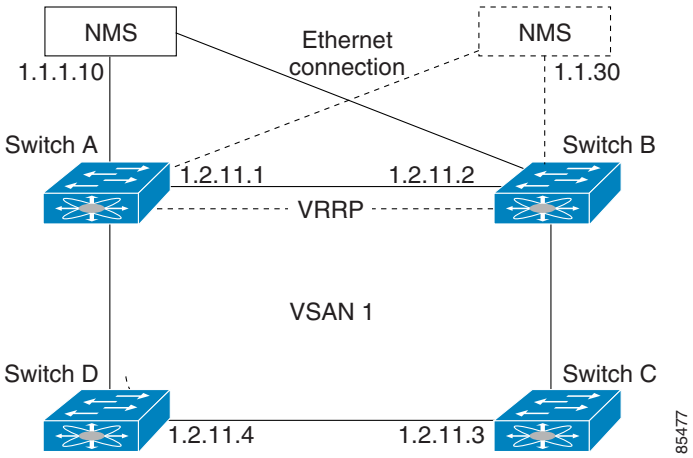

In Figure 22-7, the fabric example has two virtual router groups (VR1 and VR 2) because a virtual router cannot span across different types of interfaces. In both switch 1 and switch 2, the Ethernet interface is in VR 1 and the FC interface is in VR 2. Each virtual router is uniquely identified by the VSAN interface and the VR ID.

Figure 22-7 Redundant Gateway

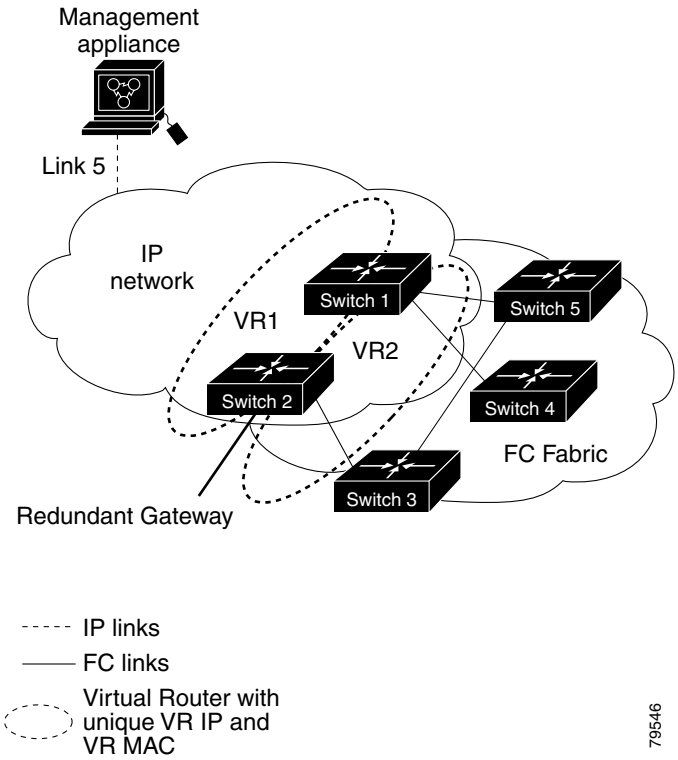

Creating or Removing a Virtual Router

All VRRP configurations should be replicated across switches in a fabric that runs VRRP.

Enabling a Virtual Router

By default, a virtual router is always disabled (**shutdown**). VRRP can be configured only if this state is disabled. Be sure to configure at least one IP address before attempting to enable a VR.

Adding an IP Address for a Virtual Router

One primary IP address and multiple secondary addresses can be configured for a switch. If the configured IP address is the same as the interface IP address, this switch automatically owns the IP address.

Viewing IP Address Information

To view IP addresses of the switches in the current fabric from the Fabric Manager, choose Switches from the menu tree.

The Information pane displays IP address information for multiple switches.

Managing IP Addresses for VRRP

To manage IP addresses for virtual routers from Device Manager, follow these steps:

Step 1 Choose **IP > VRRP**. You see the **Operations** tab of the VRRP dialog box.

Step 2 Click the **IP Addresses** tab on the VRRP dialog box.

Step 3 To create a new VRRP entry, click **Create**. You see the Create VRRP IP Addresses window.

Step 4 Complete the fields in this window to create a new VRRP IP Address, and click **OK** or **Apply**.

Setting 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 a switch with the primary IP address.

Setting the Time Interval for the Advertisement Packet

The valid time range for an advertisement packet is between 1 and 255 seconds with the default being 1 (one) second. If the switch has the primary IP address, this time must be specified.

Preempting the Master Virtual Router

By default, the preempt option is enabled. An owner with priority 255 cannot be preempted. If two priorities match, the owner with the highest priority preempts the master virtual router.

The VRRP **preempt** option is not supported on IP storage Gigabit Ethernet interfaces. However, if the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.

Configuring Authentication for the Virtual Router

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.

All VRRP configurations must be duplicated

Setting the Priority Based on Interface State

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.

Configuring VRRP Operations Attributes

To configure VRRP operations attributes from Device Manager, follow these steps:

Step 1 Choose **IP > VRRP**. You see the **Operations** tab of the VRRP dialog box.

Step 2 Configure operations attributes for the virtual router.

Step 3 To create a new VRRP entry, click **Create**. You see the Create VRRP Entry window.

Step 4 Complete the fields in this window to create a new VRRP entry, and click **OK** or **Apply**.

Default Settings

Table 22-3 lists the default settings for IP features.

Table 22-3 Default IPFC Settings

Parameters	Default
VSAN IP interface configuration	No IP address is assigned by default.
IP routing	Disabled.
Domain lookup	Disabled.
Domain name	Enabled.
Domain list	No domains are configured.
Name server	No servers are configured.
Virtual router	Disabled (shutdown).
Virtual router priority for switches with secondary IP address	100.
Virtual router priority for switches with primary IP address	255.
Time interval between advertisement frames	1 second.
Preempting master VR	Enabled.
VRRP security authentication	No authentication.
Interface state tracking	Disabled.

Note ICMP redirect packets are always rejected.

Enabling or Disabling IP Forwarding

To view or change the IP forwarding configuration of the switches in the current fabric, perform the following steps.

Step 1 Choose **IP > Forwarding** from the Fabric Manager menu tree.

Step 2 To enable IP forwarding for a specific switch, check the **RoutingEnabled** check box.

Viewing Information and Statistics

You can monitor a variety of information and statistics about IP services from Fabric Manager and Device Manager.

This section includes the following topics:

- Viewing VRRP Statistics, page 22-15
- Viewing TCP Information and Statistics, page 22-15
- Viewing UDP Information and Statistics, page 22-15
- Viewing IP Statistics, page 22-16
- Viewing ICMP Statistics, page 22-16

Viewing VRRP Statistics

To monitor VRRP statistics, click the **Statistics** tab on the VRRP dialog box. The VRRP dialog box with the Statistics tab selected is displayed.

Viewing TCP Information and Statistics

To view TCP information from the Device Manager, choose **Mgmt TCP/UDP** from the IP menu.

To monitor TCP statistics from the Fabric Manager, choose **IP > Mgmt Statistics** from the menu tree and click the **TCP** tab. To monitor TCP statistics from the Device Manager, choose **Statistics** from the IP menu and view the **TCP** tab.

Viewing UDP Information and Statistics

To view User Datagram Protocol (UDP) information, from the Device Manager, choose **Mgmt TCP/UDP** from the IP menu and click the **UDP** tab.

To monitor UDP traffic from the Fabric Manager, choose **IP > Mgmt Statistics** from the menu tree and click the **UDP** tab. From Device Manager, choose **Statistics** from the IP menu and click the **UDP** tab.

The Fabric Manager Information pane displays TCP traffic information for multiple switches. The Device Manager dialog box displays information for a single switch.

Viewing IP Statistics

To monitor IP statistics from the Fabric Manager, choose **IP >Mgmt Statistics** from the menu tree and click the **IP** tab. From Device Manager, select **Statistics** from the IP menu and click the **IP** tab.

The Fabric Manager Information pane displays IP statistics for multiple switches. The Device Manager dialog box displays information for a single switch.

Viewing ICMP Statistics

To monitor statistics for ICMP packets received, select **IP > Mgmt Statistics** from the menu tree and click the **ICMP In** tab. To monitor statistics for ICMP packets transmitted from the Fabric Manager, select **IP > Mgmt Statistics** from the menu tree and click the **ICMP Out** tab.

To monitor ICMP statistics from Device Manager, select **Statistics** from the IP menu and click the **ICMP** tab.

The Fabric Manager Information pane displays information for multiple switches. The Device Manager dialog box displays information for a single switch.

In the Device Manager, a prefix (In or Out) identifies whether the packets are received or transmitted. In the Fabric Manager, separate tabs on the Information pane are provided for incoming and outbound ICMP traffic and this prefix is omitted.

Configuring FICON

FI-bre CON-nection (FICON) interface capabilities enhances 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 inband management of the switch from FICON processors.

FICON features can be implemented in any switch in the Cisco MDS 9000 Family running SAN-OS Release 1.3(x) or above. While no hardware changes are required, you do need the MAINFRAME_PKG license to configure FICON parameters.

This chapter contains the following topics:

- About FICON, page 23-2
- MDS-Specific FICON Advantages, page 23-3
- FICON Port Numbering, page 23-6
- MDS FICON Prerequisites, page 23-11
- Enabling FICON, page 23-12
- Creating FICON VSANs (enabling FICON) Using Fabric Manager, page 23-12
- Creating FICON VSANs (enabling FICON) Using Device Manager, page 23-13
- Deleting FICON VSANs (Disabling FICON), page 23-13
- Viewing FICON Director History, page 23-14
- Configuring Code Page, page 23-14
- Configuring the FC ID Last Byte, page 23-14
- Automatically Saving the Running Configuration, page 23-14
- Binding Port Numbers to PortChannels, page 23-15
- Binding Port Numbers to FCIP Interfaces, page 23-15
- Configuring FICON Ports, page 23-15
- Entering FICON Port Configuration Information, page 23-16
- Viewing FICON Port Attributes, page 23-17
- FICON Configuration Files, page 23-17
- Creating FICON Files, page 23-18
- Deleting FICON Files, page 23-19
- Copying FICON Files, page 23-19

- Port Swapping, page 23-19
- Swapping FICON Ports, page 23-20
- Clearing FICON Device Allegiance, page 23-21
- CUP Inband Management, page 23-21
- Configuring Fabric Binding, page 23-21
- Activating Fabric Binding, page 23-24
- Deactivating Fabric Binding, page 23-24
- Fabric Binding CopyActive to Config, page 23-24
- Creating a Fabric Binding Configuration, page 23-25
- Deleting a Fabric Binding Configuration, page 23-25
- Viewing Fabric Binding Active Database, page 23-25
- Viewing Fabric Binding Violations, page 23-25
- Clearing Fabric Binding Statistics, page 23-26
- Viewing EFMD Statistics, page 23-26
- Displaying RLIR Information, page 23-26

About FICON

The Cisco MDS 9000 Family supports Fibre Channel protocol (FCP), FICON, iSCSI, and 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.

Figure 23-1 Shared System Storage Network

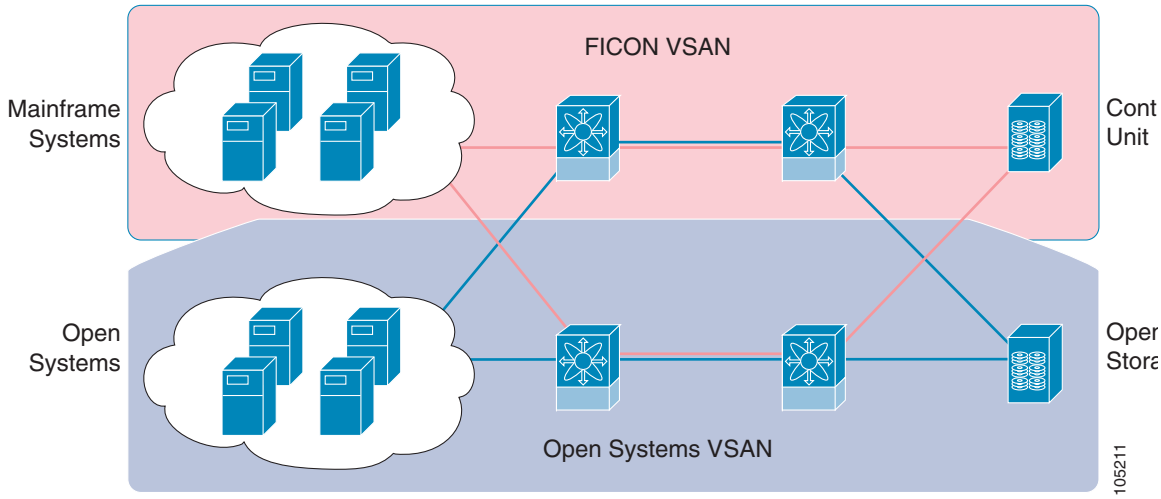

FCP and FICON are different FC4 protocols and their traffic are independent of each other. If required, devices using these protocols can be isolated using VSANs.

MDS-Specific FICON Advantages

This section explains the additional FICON advantages in Cisco MDS switches:

This section contains the following topics:

- Fabric-Optimization with VSANs, page 23-3
- FCIP Support, page 23-4
- PortChannel Support, page 23-5
- VSANs for FICON and FCP Intermixing, page 23-5
- MDS-Supported FICON Features, page 23-5

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 provides a common redundant physical infrastructure.

Figure 23-2 VSAN-specific Fabric Optimization

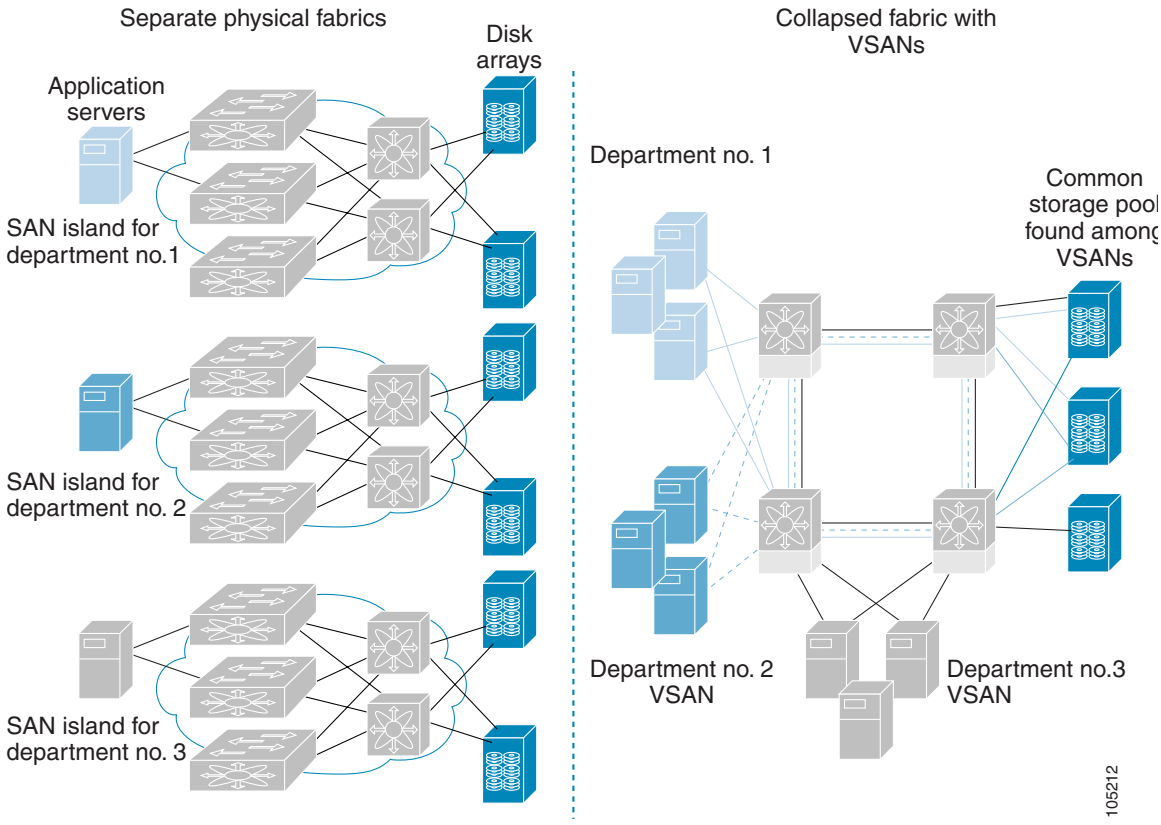

VSANs enable global SAN consolidation by allowing you to convert existing SAN islands into virtual SAN islands on a single physical network. It provides hardware-enforced security and separation between applications or departments to allow coexistence on a single network. It also allows virtual rewiring to consolidate your storage infrastructure. You can move assets between departments or applications without the expense and disruption of physical relocation of equipment.

While you can configure up to 256 VSANs in any Cisco MDS switch, you can enable FICON in only 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 MDS 9216 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 MDS 9000 platform, storage replication services such as IBM PPRC and XRC can be extended over metro to global distances using ubiquitous IP infrastructure and simplifying business continuance strategies.

The Cisco MDS implementation of FICON provides support for IP tunneling to efficiently consolidate SANs over WAN distances. IP tunnels enable a globally accessible storage infrastructure. Refer to the for further information on FCIP.

PortChannel Support

The Cisco MDS implementation of FICON provides support for efficient utilization and increased availability of inter-switch links necessary to build stable large-scale SAN environments. PortChannels ensure an enhanced ISL availability and performance in Cisco MDS switches.

VSANs for FICON and FCP Intermixing

Cisco MDS 9000 Family FICON-enabled switches simplify deployment of even the most complex intermix 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 intermix 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 SAN-OS software. The challenge of mixing Fibre Channel Protocol (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 intermixing at the port level. If these protocols are intermixed in the same switch, you can use VSANs to isolate FCP and FICON ports.

Tip When creating an intermix 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.

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 all Cisco MDS 9500 Series as well as the Cisco MDS 9216 Switch. (refer to the Cisco MDS 9500Series and the Cisco MDS 9216 Switch Hardware Installation Guides).
- 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 2/1-Gbps, autosensing FICON or Fibre Channel FCP ports in any combination in a single chassis and up to 768 Fibre Channel ports in a single rack--1.44 Tbps of internal system bandwidth ensures smooth integration of future 10-Gbps modules.
- Infrastructure protection--Common software releases infrastructure protection is available across all Cisco MDS 9000 platforms.
- 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 intermix support.

- Port-level configurations:
  - BB_credits for each port.
  - Port security for each port.
  - Enable beaconing for ports and the director unit
  - Configure an alias name, instead of the WWN, for switches and attached node devices.
  - 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), Fibre Channel Security Protocol (FC-SP), VSANs, hardware-enforced zoning, LUN zoning, read-only zones, ACLs, fabric binding and VSAN-based access control.
  - View the local accounting log to locate FICON events.
  - 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.
  - Port address-based configurations--port name, blocked or unblocked state, and the prohibit connectivity attributes.
  - 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.
- Store and apply configuration files.
- FICON and Open Systems Management Server features if installed.
- Enhanced Cascading Support
- Set the date and time on the switch.
- Configure SNMP trap recipients and community names.
- Call Home configurations--director name, location, description, and contact person.
- Configure preferred domain ID, FC ID persistence, and principle switch priority.
- 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.
- Configure R_A_TOV, E_D_TOV
- Perform maintenance tasks for the director including maintaining firmware levels, accessing the director logs, and collecting data to support failure analysis.
- Display and clear port-level incident alerts

FICON Port Numbering

With reference to the FICON feature, ports in Cisco MDS switches are identified by a statically defined 8-bit value known as the port number. Port numbers are assigned based on the module and the slot in the chassis. Port numbers cannot be changed and the first port in a switch always starts with a 0.

Figure 23-3 Port Number in the Cisco MDS 9000 Family

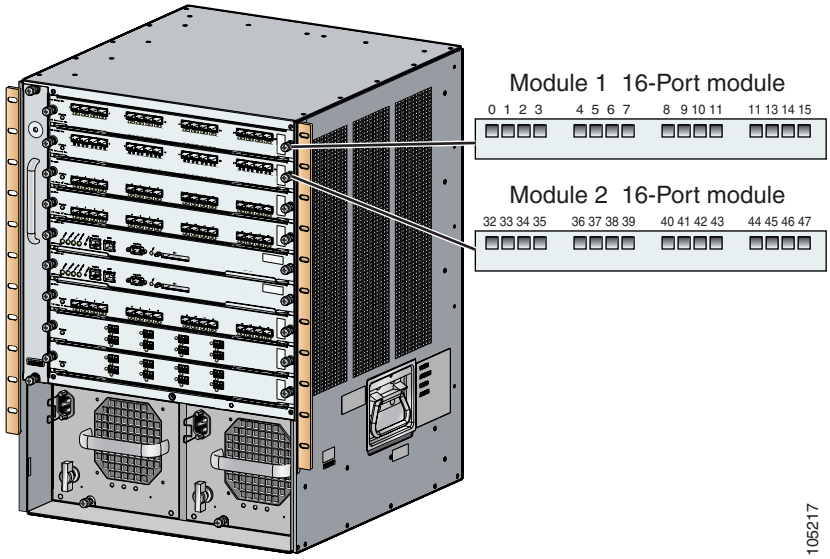

The FICON port number is assigned based on the front panel location of the port and is specific to the slot in which the module resides. Even if the module is a 16-port module, 32-port numbers are assigned to that module--regardless of the module type (16-port or 32-port), the module's physical presence in the chassis, or the port status (up or down).

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 23-1 lists the port number assignment for the Cisco MDS 9000 Family of switches and directors.

Table 23-1 Cisco MDS 9000 Family Port Number Assignments

Product	Slot Number	Implemented Port Allocation	Unimplemented Ports	Notes
To Ports	**To PortChannel/FCIP**			
Cisco MDS 9120 Switch	Not applicable	Ports 0 through 19	20 through 55	Ports 56 through 253 and Port 255
Cisco MDS 9140 Switch	Not applicable	Ports 0 through 39	40 through 65	Ports 66 through 253 and Port 255
Cisco MDS 9216 Switch	Slot 1	Ports 0 through 31	64 through 89	Ports 90 through 253 and Port 255
Slot 2	Ports 32 through 63	The first 16 port numbers in a 16-port module are used and the rest remain unused.		
Cisco MDS 9506 Director	Slot 1	Ports 0 through 31	Ports 128 through 143	Ports 144 through 253 and Port 255
Slot 2	Ports 32 through 63			

Table 23-1 Cisco MDS 9000 Family Port Number Assignments (continued)

Slot 3	Ports 64 through 95			
Slot 4	Ports 96 through 127			
Slot 5	None	Supervisor modules are not allocated port numbers.		
Slot 6	None			
Cisco MDS 9509 Director	Slot 1	Ports 0 through 31	Ports 224 through 239	Ports 240 through 253 and Port 255
Slot 2	Ports 32 through 63			
Slot 3	Ports 64 through 95			
Slot 4	Ports 96 through 127			
Slot 5	None	Supervisor module are not allocated port numbers.		
Slot 6	None			
Slot 7	Ports 128 through 159	The first 16 port numbers in a 16-port module are used and the rest remain unused.		
Slot 8	Ports 160 through 191			
Slot 9	Ports 192 through 223			

Port Addresses

By default, port numbers are the same as port addresses.

Implemented and Unimplemented Port Addresses

An implemented port refers to any port address that is available in the chassis. These numbers are identified in the Implemented Port Allocation column.

An unimplemented port refers to any port address that is not available in the chassis. These numbers are identified in the Unimplemented Ports column.

Tip An unimplemented port is prohibited from communicating with an implemented port in a FICON setup and cannot be configured.

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, port numbers 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, port numbers 48 to 63 are considered uninstalled.
- The port is not in a FICON-enabled VSAN, for example, if port number 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 (see Implemented Port Allocation To PortChannel/FCIP), for example, if interface fc 1/1 is part of PortChanne1 5, port address 0 is uninstalled in all FICON VSANs.

For example:

- If module 1 is not physically present in slot 1 in a Cisco MDS 9509 Director, port numbers 0 to 31 are considered uninstalled.
- If a 16-port module is inserted in slot 2 in a Cisco MDS 9509 Director, port numbers 38 to 63 are considered uninstalled.
- If port number 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.
- If interface fc1/1 (port address = 0) is a TE port and is configured in VSANs 1 through 20--but only VSANs 2 and 3 are FICON-enabled, then port address 0 is only installed in VSAN 2 and VSAN 3.

FCIP Port Number

You must explicitly configure FCIP port numbers. The port address for FCIP ports are configured to the range of numbers that you can use are restricted to the port numbers available in the IPS modules slot. If an IPS module is in Slot 9 in a Cisco MDS 9509 Director, the available range of port numbers is 192 through 223. The FCIP interface can be assigned any port number that is available within that range.

For example, if the FCIP port is bound to Gigabit Ethernet interface 9/1, the assigned FCIP port numbers can be 192, 193, 194, 195, 208, 209, 210, or 211.

Figure 23-4 FCIP Port Numbers in the Cisco MDS 9000 Family

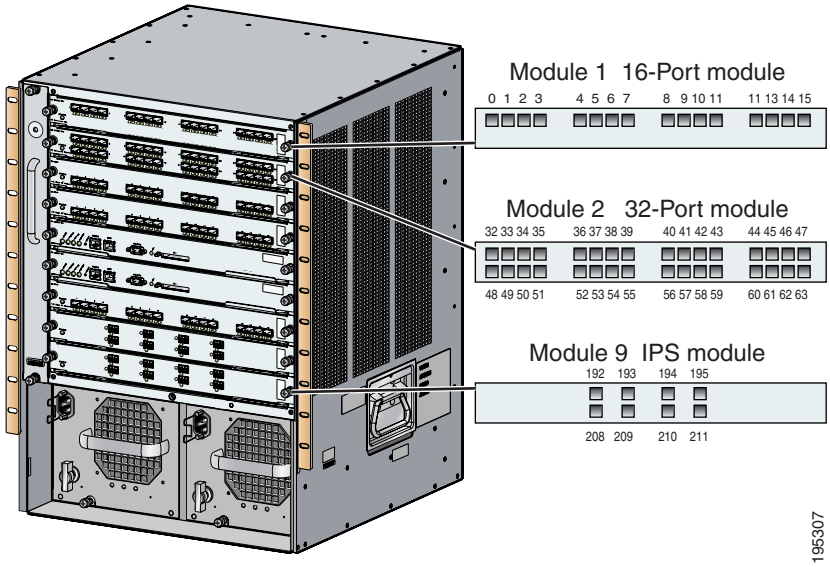

Gigabit Ethernet ports do not have a corresponding mapping to the FICON port number concept.

Port Numbering Summary

The following guidelines apply to FICON port numbers:

- Supervisor modules do not have port number assignments.
- Port number are VSAN independent--Fibre Channel port numbers 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 number is not assigned for PortChannels or for FCIP tunnels, the associated ports will not come up (see the "FCIP and PortChannel Port Numbers" section on page 23-10).

FCIP and PortChannel Port Numbers

FCIP and PortChannels cannot be used in a FICON enabled VSAN unless they are explicitly bound to a port number. Refer to"Binding Port Numbers to PortChannels" section on page 23-15 or "Binding Port Numbers to FCIP Interfaces" section on page 23-15 for configuration details.

Tip You can assign port numbers which are not in the PortChannel range and which are not bound to PortChannel or FCIP interfaces.

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.

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.

Figure 23-5 Static FC ID Allocation for FICON

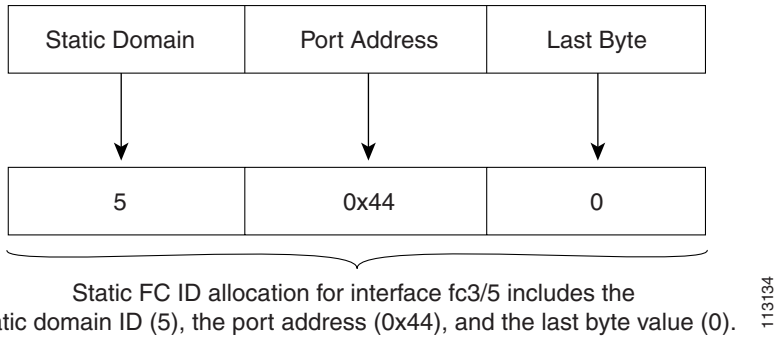

FICON Cascading

The SAN-OS software allows multiple switches in a FICON network. To configure multiple switches, you must enabled and configure fabric binding in that switch.

MDS FICON 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.
- Enable in-order delivery on the VSAN.
- Enable (and if required, configure) fabric binding on the VSAN.
- Verify that conflicting persistent FC IDs do not exist in the switch.
- Verify that the configured domain ID and requested domain ID match.
- Add the CUP (area FE) to the zone, if you are using zoning.

If any of these requirements are not met, the FICON feature cannot be enabled.

Enabling FICON

By default FICON is disabled in all switches in the Cisco MDS 9000 Family. You can enabled FICON on a per VSAN basis in one of two ways:

- Manually addressing each prerequisite.
- By using the Device Manager (refer to the Cisco MDS 9000 Family Fabric Manager User Guide for further information).

Effects of Enabling FICON

When you enable the FICON feature in Cisco MDS switches, the following apply:

- The IPL configuration file is automatically created.
- You cannot disable in-order delivery, fabric binding, or static domain ID configurations.
- 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.

Creating FICON VSANs (enabling FICON) Using Fabric Manager

The VSAN that is created here does not need to be a new VSAN. It is a new FICON VSAN. When a new FICON VSAN is created, static (insistent) domain IDs, in-order delivery, and fabric binding must be enabled so the FICON VSAN can operate. When you enable the FICON feature in Cisco MDS switches, the following apply:

- The IPL configuration file is automatically created (see the “FICON Configuration Files” section).
- You cannot disable in-order delivery, fabric binding, or static (insistent) domain ID configurations.

If you specify an existing VSAN with operational traffic to be used for the FICON VSAN, the traffic will be disrupted. In this case, a warning message is displayed before you create the FICON VSAN.

To create a FICON VSAN using Fabric Manager, follow these steps:

Step 1 From Fabric Manager, right click on **All VSANs** in the Logical pane, and click **Create VSAN**.

You see the Create VSAN dialog box.

Step 2 Choose 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 admin state you would like.

Step 6 If this is to be a FICON VSAN, check the **FICON** check box.

Step 7 To enable fabric binding for the selected switches, check that check box.

Step 8 Click **Create** to create the new VSAN, or click **Close** to close the dialog box without creating the VSAN.

Step 9 Open Device Manager for each switch in the VSAN.

Step 10 Choose **FICON > VSANs**. You see the FICON VSANs dialog box.

Step 11 Enter the VSAN membership information.

Creating FICON VSANs (enabling FICON) Using Device Manager

The VSAN that is created here does not need to be a new VSAN. It is a new FICON VSAN. When a new FICON VSAN is created, static (insistent) domain IDs, in-order delivery, and fabric binding must be enabled so the FICON VSAN can operate. When you enable the FICON feature in Cisco MDS switches, the following apply:

- The IPL configuration file is automatically created (see the "FICON Configuration Files" section on page 23-17).
- You cannot disable in-order delivery, fabric binding, or static (insistent) domain ID configurations.

If you specify an existing VSAN with operational traffic to be used for the FICON VSAN, the traffic will be disrupted. In this case, a warning message is displayed before you create the FICON VSAN.

To create a FICON VSAN, follow these steps:

Step 1 From Device Manager, choose **FICON > VSANs**. You see the FICON VSANs/Files configuration dialog box.

Step 2 Ensure that the VSANs tab is enabled.

Step 3 Click **Create**. You see the Create FICON VSANs dialog box.

Step 4 Enter the VSAN ID.

Step 5 Enter the Domain ID.

Step 6 Click **Create** to create the new VSAN, or click **Close** to close the dialog box without creating the VSAN.

Deleting FICON VSANs (Disabling FICON)

To delete a FICON VSAN, follow these steps:

Step 1 From Device Manager, choose **FICON > VSANs**. You see the FICON VSAN configuration dialog box.

Step 2 Ensure that the VSANs tab is enabled.

Step 3 Click anywhere in the row for the VSAN you want to delete.

Step 4 Click **Delete** to delete the VSAN.

Note Deleting the VSAN will also delete the associated FICON configuration file, and the file cannot be recovered.

Viewing FICON Director History

To view FICON director history, follow these steps:

Step 1 From Device Manager, choose **FICON > VSANs**. You see the FICON VSAN configuration dialog box.

Step 2 Ensure that the VSANs tab is enabled.

Step 3 Click anywhere in the row for the VSAN that you are interested in.

Step 4 Click **Director History** to display a history of FICON-related changes to this switch.

Configuring Code Page

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 FC ID Last Byte

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.

Automatically Saving the Running Configuration

When active equals saved is enabled in a VSAN, any FICON changes to the block, prohibit or port address name are immediately written to the IPL file (see the "FICON Configuration Files" section on page 23-17).

If active equals saved is enabled in any FICON-enabled VSAN in the fabric, all non-FICON configuration changes are automatically saved to persistent storage (implicit copy running start) when a configuration is changed. In this case, the current active configuration across all VSANs is saved to persistent storage.

The table below displays the results of active equals saved being enabled in at least one FICON VSAN.

Note If your switch or fabric consists of multiple FICON -enabled VSANs, and one of these VSANs have active equals saved enabled, changes made to the non-FICON configuration results in all configurations being saved to the startup configuration.

Binding Port Numbers to PortChannels

You can bind (or associate) a PortChannel with a FICON port number to bring up that interface.

Binding Port Numbers to FCIP Interfaces

You can bind (or associate) an FCIP interface with a FICON port number to the selected PortChannel interface.

Configuring 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.

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 in order 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.

Blocking Ports

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 is not allowed on that port.

Physical Fibre Channel port blocks will continue to transmit Off-Line State (OLS) primitive sequence on a blocked port.

Note The shutdown/no shutdown port state is independent of the block/no block port state. If a port is shutdown, unblocking that port will not initialize the port.

Caution You cannot block or prohibit CUP port (0XFE).

Prohibiting Ports

To prevent implemented ports (see the"Implemented and Unimplemented Port Addresses" section on page 23-9) from talking to each other, you can configure two or more ports to be in a prohibited state. If you prohibit ports, the specified ports are prevented from communicating with each other.

Note Unimplemented ports are always prohibited.

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 nor in TE mode.

Entering FICON Port Configuration Information

Note To view the latest FICON information, you must click the Refresh button. See the "FICON Information Refresh Note" section on page 23-15 for more information.

To display FICON Port Configuration information, follow these steps:

Step 1 From Device Manager, choose **FICON > VSANs**. You see the FICON VSAN configuration dialog box.

Step 2 Ensure that the VSANs tab is enabled.

Step 3 Click anywhere in the row for the VSAN for which you want to configure port information.

Step 4 Click **Port Configuration** to display the Port Configuration dialog box.

Step 5 Enter the port configuration information. Click **Apply** to save the configuration information, or click **Cancel** to close the dialog box without saving the changes.

Viewing FICON Port Attributes

Note To view the latest FICON information, you must click Refresh. See the "FICON Information Refresh Note" section on page 23-15 for more information.

To view FICON port attributes, follow these steps:

Step 1 From Device Manager, choose **FICON > VSANs**. You see the FICON VSAN configuration dialog box.

Step 2 Ensure that the VSANs tab is enabled.

Step 3 Click anywhere in the row for the VSAN that you are interested in.

Step 4 Click **Port Attributes** to display the Port Attributes 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 TM. These files can be read and written by IBM hosts using the inband CUP protocol. Additionally, you can use the Cisco MDS CLI or FM applications to operate these FICON configuration files

Multiple FICON configuration files with the same name can exist in the same switch, provide 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 uses the startup FICON configuration file, called IPL. This file is created with a default configuration as soon as FICON is enabled on a VSAN.

FICON configuration files contain the following configuration for each implemented port address:

- Block
- Prohibit mask
- Port address name

Refer to the

for details on the normal configuration files used by Cisco MDS switches. This configuration file includes FICON enabled attribute for a VSAN, port number mapping for port channels and FCIP interfaces, port number to port address mapping, port and trunk allowed VSAN configuration for ports, in-order guarantee, configuring static domain ID, and fabric binding configuration.

Caution When FICON is disabled on a VSAN, all the FICON configuration files are irretrievably lost.

Accessing FICON Configuration Files

Only one user can access the configuration file at any given time:

- While this file is being accessed by user 1, user 2 cannot access this file.
- When user 2 does attempt to access this file, an error is issued to user 2.
- If user 1 has been 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 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.

Editing FICON Configuration Files

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 8 alphanumeric characters.

Creating FICON Files

If a specified file does not exist, it is created. Up to 16 files can be saved. Each file name is restricted to 8 alphanumeric characters.

Note To view the latest FICON information, you must click the Refresh button. See the "FICON Information Refresh Note" section on page 23-15 for more information.

To create a FICON file, follow these steps:

Step 1 From Device Manager, choose **FICON > VSANs**. You see the FICON VSANs/Files dialog box.

Step 2 Click the **Files** tab.

Step 3 Click **Create**. You see the Create FICON VSANs Files dialog box.

Step 4 Enter the VSAN ID.

Step 5 Enter the File Name.

Step 6 Enter the Description.

Step 7 Click **Create** to create the new file, or click **Close** to close the dialog box without creating the file.

Deleting FICON Files

Note To view the latest FICON information, you must click Refresh. See the "FICON Information Refresh Note" section on page 23-15 for more information.

To delete a FICON file, follow these steps:

Step 1 From Device Manager, choose **FICON > VSANs**. You see the FICON VSANs/Files dialog box.

Step 2 Click the **Files** tab.

Step 3 Click anywhere in the row for the file you want to delete.

Step 4 Click **Delete** to delete the file.

Copying FICON Files

The SAN-OS software maintains different configuration files to support a FICON network. These configuration files can be saved using Device Manager. FICON configuration files do not contain the following information that is normally saved with the running configuration:

Note To view the latest FICON information, you must click the Refresh button. See the "FICON Information Refresh Note" section on page 23-15 for more information.

- Port number to port address mapping
- PortChannel to port number mapping
- Port swap occurrences
- FICON enabled VSANs

FICON configuration files are independent of these parameters. Instead, this information is stored in persistent storage as they can be modified independent of the startup configuration.

To copy a FICON file, follow these steps:

Step 1 From Device Manager, choose **FICON > VSANs**. You see the FICON VSANs/Files dialog box.

Step 2 Click the **Files** tab.

Step 3 Click to highlight the row for the file you want to copy.

Port Swapping

The port swap FICON feature is only provided for maintenance purposes and is supported in all switches in the Cisco MDS 9000 Family support this feature.

Swapping the ports causes all configuration associated with old-port-number and new port-number to be swapped, including VSAN configurations. This command is only associated with the two ports in concerned. You must issue this VSAN-independent command from the EXEC mode.

Tip If the **active equals saved** option is enabled on any FICON VSAN, then the swapped configuration is automatically saved to startup. Otherwise, you must explicitly copy the running startup immediately after swapping the ports.

MDS switches also allow port swapping for non-existent ports as specified below:

- Only FICON-specific configurations (prohibit, block, and port address mapping) is swapped.
- No other system configuration is swapped.
- All other system configurations are only maintained for existing ports.

Once you swap ports, the switch automatically performs the following actions:

- Shuts down both the old and new ports
- Swaps the port configuration
- Any attempt to bring the port up using the after swap noshut option after the new-port-number will not work unless you explicitly issue the **no shutdown** command to resume traffic.

Port Swapping Guidelines

Be sure to follow these guidelines when using the FICON port swap 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 SAN-OS software performs a compatibility check. If the two ports have incompatible configuration, 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 swap operation is rejected.
- If ports have default values (for some incompatible parameters), then port swap is allowed to go through and the ports retain their default values. If you swap a 16-port module with a 32-port module, the BB_credits will no longer be compatible and the ports can be swapped. If BB_credits are not configured, the default settings will still be in effect at the time of the swap.

The 32-port module guidelines also apply for port swapping configurations.

Swapping FICON Ports

Note To view the latest FICON information, you must click the Refresh button. See the "FICON Information Refresh Note" section on page 23-15 for more information.

To swap ports, follow these steps:

Step 1 Choose two Fibre Channel ports by holding down the **CTRL** key and clicking on them.

Step 2 Choose **FICON > Swap Selected Ports**.

Clearing FICON Device Allegiance

FICON requires serialization of access between 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.

CUP Inband 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 inband management using the IBM S/A OS/390 I/O operations console.

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 MDS switches.

Host communication includes control functions like blocking/unblocking ports, as well as monitoring and error reporting functions.

Configuring Fabric Binding

The SAN-OS 1.3(x) 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 contains the following topics:

- Port Security versus Fabric Binding, page 23-22
- Enforcing Fabric Binding, page 23-22
- Enabling Fabric Binding, page 23-23
- Configuring a List of sWWNs, page 23-23
- Activating Fabric Binding, page 23-23

Port Security versus Fabric Binding

Port security and fabric binding are two independent features that can be configured to complement each other. (See Table 23-2,)

Table 23-2 Port Security and Fabric Binding Descriptions

Fabric Binding	Port Security
Configured using a set of sWWN and a persistent Domain ID.	Configured using pWWNs/nWWNs or fWWNs/switch WWNs.
Binds the fabric at the switch level.	Binds devices at the interface level.
Only the configured sWWN stored in the fabric binding database will be authorized to participate in the fabric.	Allows a preconfigured set of Fibre Channel devices to logically connect to a SAN port(s). The switchport, 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 (list).
Activation is required on a per VSAN basis.	Activation is required on a per VSAN basis.
User defines specific switches which are allowed to connect to the fabric, regardless of the physical port to which the peer switch is connected.	User specifies the specific physical port(s) to which another device can connect
Does not learn logging in switches.	Learns about switches/devices if in learning mode.

Port-level Checking for xE ports

- Switch login—Uses both Port Binding as well as the Fabric Binding feature for a given VSAN.
- Binding checks are done on the port VSAN:
  - E-port security binding check done on port VSAN.
  - TE-port security bindings check done in each VSAN allowed.

While port security complements fabric binding, they are independent features and can be enabled or disabled separately.

Enforcing Fabric Binding

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. However enforcement of fabric binding at the time of activation happens only if the VSAN is a FICON VSAN.

The fabric binding feature requires all sWWNs connecting to a switch and their persistent domain IDs to be part of the fabric-binding active database.

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 participate 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.

Configuring a List of sWWNs

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.

Activating Fabric Binding

The fabric binding maintains a configuration database (config-database) and an active database. The config-database is a read-write database which 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 login.

By default, the fabric binding feature is not activated. You cannot activate the switch if entries existing in the config 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 the **force** option to override these situations.

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 due to fabric binding restrictions will be reinitialized.

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.

Activating Fabric Binding

The fabric binding feature must be enabled in each switch in the fabric that participate in the fabric binding. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

To activate fabric binding, follow these steps:

Step 1 From Device Manager, choose **FICON > Fabric Binding**. You see the Fabric Binding dialog box.

Step 2 Ensure that the Actions tab is enabled.

Step 3 Click in the Actions column for the VSAN(s) for which you want to activate fabric binding.

Step 4 Choose **Activate** or **Force Activate**.

Step 5 Click **Apply** to activate the fabric binding, or click **Close** to close the dialog box without activating fabric binding for the selected VSAN(s).

Deactivating Fabric Binding

To deactivate fabric binding, follow these steps:

Step 1 From Device Manager, choose **FICON > Fabric Binding**. You see the Fabric Binding dialog box.

Step 2 Ensure that the Actions tab is enabled.

Step 3 Click in the Actions column for the VSAN(s) for which you want to deactivate fabric binding.

Step 4 Choose **Deactivate**.

Step 5 Click **Apply** to deactivate the fabric binding, or click **Close** to close the dialog box without deactivating fabric binding for the selected VSAN(s).

Fabric Binding CopyActive to Config

To copy the active fabric binding to the configuration file, follow these steps:

Step 1 From Device Manager, choose **FICON > Fabric Binding**. You see the Fabric Binding dialog box.

Step 2 Ensure that the Actions tab is enabled.

Step 3 Click in the CopyActive ToConfig column for the VSAN(s) for which you want to copy fabric binding.

Step 4 Click **Apply** to copy the fabric binding, or click **Close** to close the dialog box without copying the fabric binding for the selected VSAN(s).

Creating a Fabric Binding Configuration

To create a fabric binding configuration, follow these steps:

Step 1 From Device Manager, choose **FICON > Fabric Binding**. You see the Fabric Binding dialog box.

Step 2 Click the **Config Database** tab.

Step 3 Click **Create** to display the Create Fabric Binding Config Database dialog box.

Step 4 Enter the VSAN ID, the peer WWN, and the domain ID.

Step 5 Click **Create** to create the fabric binding configuration, or click **Close** to close the dialog box without creating the fabric binding configuration.

Deleting a Fabric Binding Configuration

To delete a fabric binding configuration, follow these steps:

Step 1 From Device Manager, choose **FICON > Fabric Binding**. You see the Fabric Binding dialog box.

Step 2 Click the **Config Database** tab.

Step 3 Click in the row for the VSAN for which you want to delete the fabric binding configuration.

Step 4 Click **Delete** to delete the fabric binding configuration, or click **Close** to close the dialog box without deleting the fabric binding configuration.

Viewing Fabric Binding Active Database

To view the fabric binding active database, follow these steps:

Step 1 From Device Manager, choose **FICON > Fabric Binding**. You see the Fabric Binding dialog box.

Step 2 Click the **Active Database** tab to display the active database.

Viewing Fabric Binding Violations

To view fabric binding violations, follow these steps:

Step 1 From Device Manager, choose **FICON > Fabric Binding**. You see the Fabric Binding dialog box.

Step 2 Click the **Violations** tab to display fabric binding violations.

Clearing Fabric Binding Statistics

To clear fabric binding statistics, follow these steps:

Step 1 From Device Manager, choose **FICON > Fabric Binding**. You see the Fabric Binding dialog box.

Step 2 Click the **Statistics** tab to display the fabric binding statistics.

Step 3 Check the check box in the Clear column for the VSAN(s) for which you want to clear statistics.

Step 4 Click **Apply**.

Viewing EFMD Statistics

To view EFMD statistics, follow these steps:

Step 1 From Device Manager, choose **FICON > Fabric Binding**. You see the Fabric Binding dialog box.

Step 2 Click the **EFMD Statistics** tab to display EFMD statistics.

Displaying RLIR Information

The Registered Link Incident Report (RLIR) application provides a method for a switchport to send a LIR to a registered Nx-port.

When a Link Incident Record (LIR) is detected in FICON-enabled switches in the Cisco MDS 9000 Family form a RLIR Extended Link Service (ELS) and sends it to the members in it's Established Registration List (ERL).

In case of multi-switch topology, a Distribute Registered Link Incident Record (DRLIR) Inter Link Service (ILS) are sent to all reachable remote domains along with the RLIR ELS. On receiving the DRLIR ILS, the switch extracts the RLIR ELS and sends 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 application is highly available and the data is stored to persistent storage when the running configuration is saved.

Viewing the RLIR information shows the complete statistics of LIRR, RLIR, and DRLIR frames. It lists the number of frames received, sent, and rejected. Specify the VSAN ID for per VSAN statistics, If you do not specify the VSAN ID, then the statistics is shown for all active VSANs.

Configuring IP Storage

Cisco MDS 9000 Family IP storage (IPS) services modules 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 allows IP hosts to access Fibre Channel storage using iSCSI protocol.

FCIP and iSCSI features are specific to the IPS module and can be implemented in Cisco MDS 9216 switches or Cisco MDS 9500 Directors running Cisco MDS SAN-OS Release 1.1(x) or above.

This chapter contains the following topics:

- IP Storage Services Module, page 24-1
- Configuring Gigabit Ethernet Interfaces, page 24-2
- Configuring FCIP, page 24-7
- Configuring iSCSI, page 24-24
- Configuring Storage Name Services, page 24-42
- Default IP Storage Settings, page 24-44
- Using the IP Filter Wizard, page 24-45
- Creating IP Profiles, page 24-45
- Adding IP Filters to Profiles, page 24-46
- Associating IP Profiles to Interfaces, page 24-46
- Deleting IP Profiles, page 24-47
- Deleting IP Filters, page 24-47

IP Storage Services Module

The IPS services module (IPS module) allows you to use FCIP and iSCSI features. It integrates seamlessly into the Cisco MDS 9000 Family, and supports the full range of features available on other switching modules, including VSANs, security, and traffic management.

The following types of IPS modules are currently available for use in any Cisco MDS 9216 switch or any switch 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.

Each port in either module can be configured to support FCIP protocol, 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. The figure below depicts the FCIP scenarios in which the IPS module is used.

Figure 24-1 FCIP Scenarios

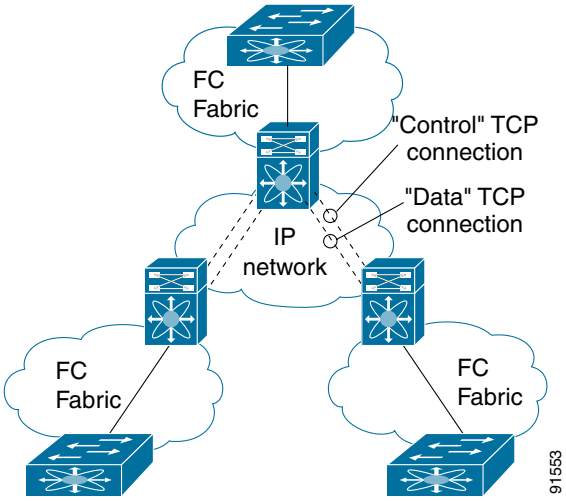

- iSCSI—The IPS module provides IP hosts access to Fibre Channel storage devices. The IP host sends SCSI commands encapsulated in iSCSI protocol data units (PDUs) to a MDS 9000 IPS port over a Transmission Control Protocol (TCP)/Internet Protocol (IP) connection. At this point, the commands are routed from an IP network into a Fibre Channel network and forwarded to the intended target. Figure 24-2 depicts the iSCSI scenarios in which the IPS module is used.

Figure 24-2 iSCSI Scenarios

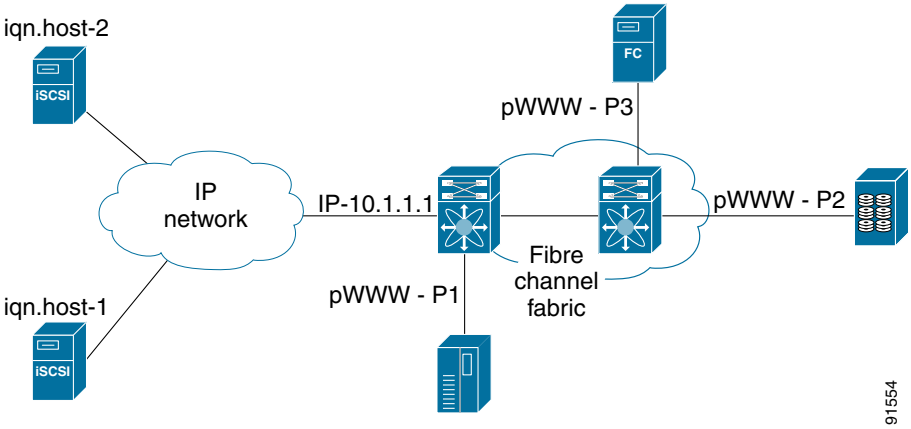

Configuring Gigabit Ethernet Interfaces

This section contains the following topics:

- About Gigabit Ethernet Interfaces, page 24-3
- Basic Gigabit Ethernet Configuration, page 24-3

- About VLANs for Gigabit Ethernet, page 24-3
- Verifying Gigabit Ethernet Connectivity, page 24-4
- Gigabit Ethernet High Availability, page 24-4
- Configuring CDP, page 24-6
- IPS Core Dumps, page 24-6

About Gigabit Ethernet Interfaces

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On the IPS 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.

A new port mode, called **IPS**, is defined for Gigabit Ethernet ports on the IPS module. IP storage ports are implicitly set to IPS mode, so they can only be used to perform iSCSI and FCIP storage functions. IP storage ports do not bridge Ethernet frames or route other IP packets.

Tip Gigabit Ethernet ports on the IPS 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 24-3 depicts a basic Gigabit Ethernet configuration.

Figure 24-3 Gigabit Ethernet Configuration

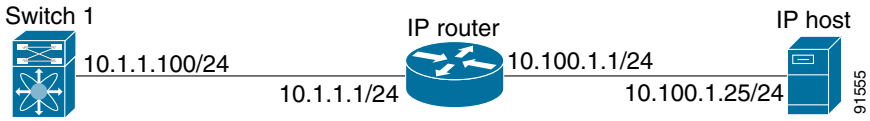

You can configure the switch to receive and transfer large (or jumbo) frames on a port. The default IP MTU frame size is 1500 bytes for all Ethernet ports. By configuring jumbo frames on a port, the MTU size can be increased to 9000 bytes. The following example sets the size to 3000 bytes. Independent of the MTU size, the IPS module does not pack multiple IP frames (converted to FCIP or to iSCSI).

The minimum MTU size for a port running iSCSI is 512 bytes.

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.

IPS 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 IPS port, configure subinterfaces—one for each VLAN. 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>).

If the IPS 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 is configured as a trunking port.
- The encapsulation is set to 802.1Q and not ISL, which is the default.

Verifying Gigabit Ethernet Connectivity

The ping command sends echo request packets out to a remote device at an IP address that you specify.

Once the Gigabit Ethernet interfaces are connected with valid IP addresses, verify the interface connectivity on each switch using the **ping** command. Ping the IP host using the IP address of the host to verify that the static IP route is configured correctly.

If the connection fails, verify the following, and repeat the ping command:

- 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 High Availability

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

Configuring VRRP

VRRP provides a redundant alternate path to the Gigabit Ethernet port for iSCSI and FCIP services.

VRRP provides IP address fail over protection to an alternate Gigabit Ethernet interface so the IP address is always available.

Figure 24-4 VRRP Scenario

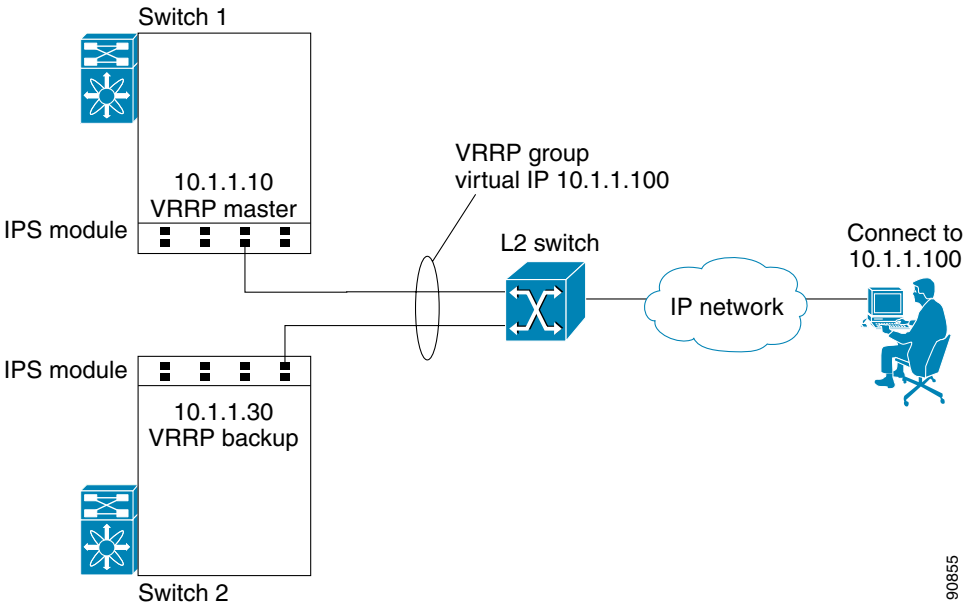

In Figure 24-4, all members of the VRRP group must be IP storage Gigabit Ethernet ports. VRRP group members can be one or more of the following interfaces:

- One or more interfaces in the same IPS module
- Interfaces across IPS modules in one switch
- Interfaces across IPS modules in different switches
- Gigabit Ethernet subinterfaces
- Ethernet PortChannels
- Subinterfaces

The VRRP **preempt** option is not supported on IP storage Gigabit Ethernet interfaces. However, if the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.

Configuring Ethernet PortChannels

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.

The data traffic from one TCP connection always travels on the same physical links. An Ethernet switch connecting to the MDS Gigabit Ethernet port can implement load balancing based on its IP address, its source-destination MAC address, or its IP and port. In iSCSI scenarios if the Ethernet switch is not capable of load-balancing based on the IP address or the IP port, multiple iSCSI initiators are required to take advantage of the Ethernet PortChannel feature.

The Cisco Ethernet switch's PortChannel should be configured as a static PortChannel, and not the default 802.3aa protocol.

Ethernet PortChannels can only aggregate two physical interfaces that are adjacent to each other on a given IPS module.

PortChannel members must be one of these combinations: ports 1-2, ports 3-4, ports 5-6, or ports 7-8.

Figure 24-5 Ethernet PortChannel Scenario

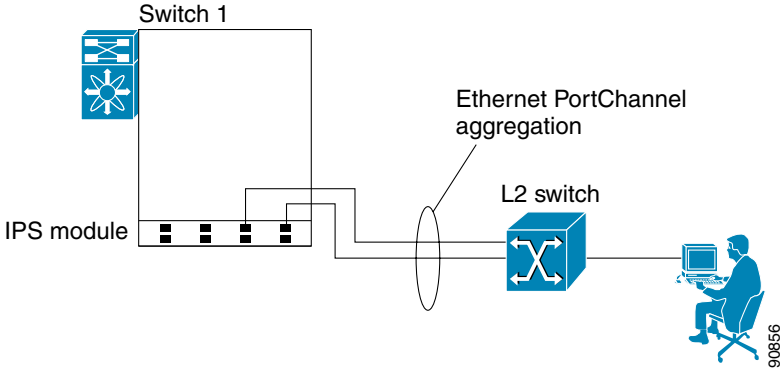

In Figure 24-5, Gigabit Ethernet ports 3 and 4 in slot 9 are aggregated into an Ethernet PortChannel.

All FCIP data traffic for one FCIP link is carried on one TCP connection. Consequently, the aggregated bandwidth will be one Gbps for that FCIP link.

PortChannel configuration specified in also apply to Ethernet PortChannel configurations.

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.

Gigabit Ethernet interfaces cannot be added to a

PortChannel if one of the following cases apply: - if the interface already has an IP address assigned, or - if subinterfaces are configured on that interface.

Configuring CDP

The Cisco Discovery Protocol (CDP) is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interface on the IPS module.

IPS Core Dumps

IPS core dumps are different from the system kernel core dumps for other modules. When the IPS module operating system (OS) unexpectedly resets, it is sometimes useful to obtain a copy of the memory image (called a IPS core dump) to identify the cause of the reset. Under that condition, the IPS module sends the core dump to the supervisor module for storage. Core dumps take up significant space and hence the level of what gets stored can be configured using one of the two options:

- Partial core dumps (default)—Each partial core dump consists of four parts (four files).
- Full core dumps—Each full core dump consists of 75 parts (75 files). This dump cannot be saved on the supervisor module due to its large space requirement.

Configuring FCIP

This section contains the following topics:

- About FCIP, page 24-7
- Enabling FCIP, page 24-9
- Basic FCIP Configuration, page 24-9
- Creating FCIP Tunnels with Device Manager, page 24-11
- Creating FCIP Tunnels with Fabric Manager, page 24-13
- Advanced FCIP Profile Configuration, page 24-13
- Advanced FCIP Interface Configuration, page 24-16
- E Port Configurations, page 24-19
- Enabling FCIP Compression, page 24-21
- Ethernet PortChannels and Fibre Channel PortChannels, page 24-24

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).

FCIP uses TCP as a network layer transport.

To configure the IPS module for FCIP, you should have a basic understanding of the following concepts:

- FCIP and VE Ports
- FCIP Links
- FCIP Profiles
- FCIP Interfaces

FCIP and VE Ports

Figure 24-6 describes the internal model of FCIP with respect to Fibre Channel inter switch links (ISLs) and Cisco's enhanced ISLs (EISLs).

FCIP defines virtual E (VE) ports, which 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 a 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.

Figure 24-6 FCIP Links and Virtual ISLs

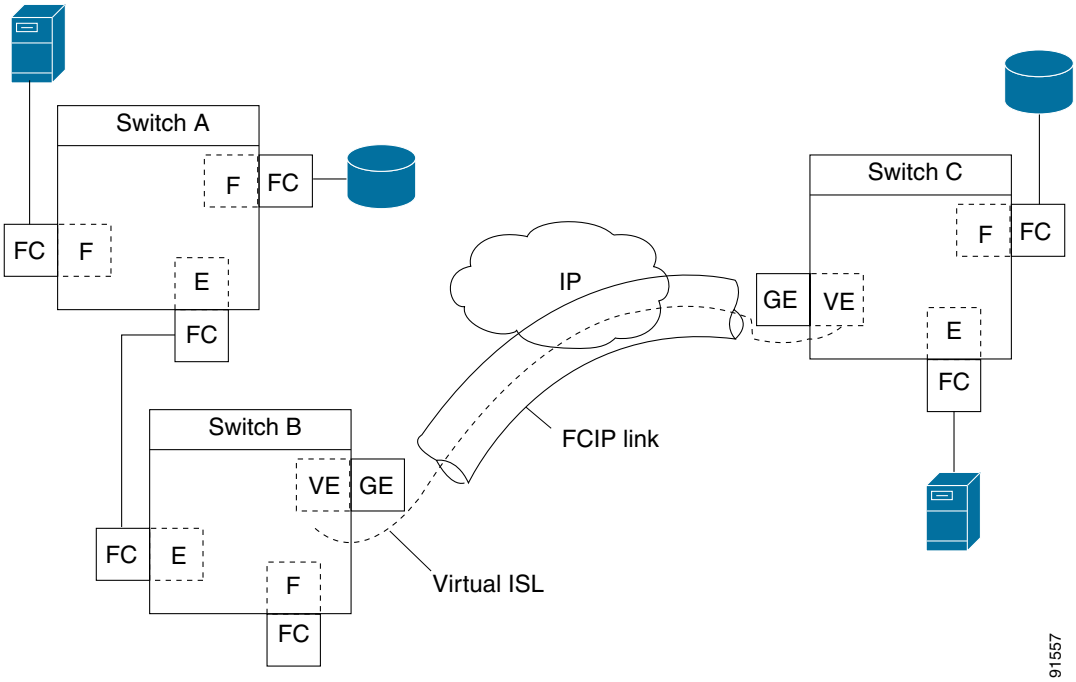

FCIP Link

FCIP links consist of one or more TCP connections between two FCIP link end points. 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 second connection is used only for Fibre Channel control frames, i.e. switch-to-switch protocol frames (all Class F) frames. This arrangement is used to provide low latency for all control frames.

To enable FCIP on the IPS module, a 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 terminates.

Figure 24-7 FCIP Profile and FCIP Links

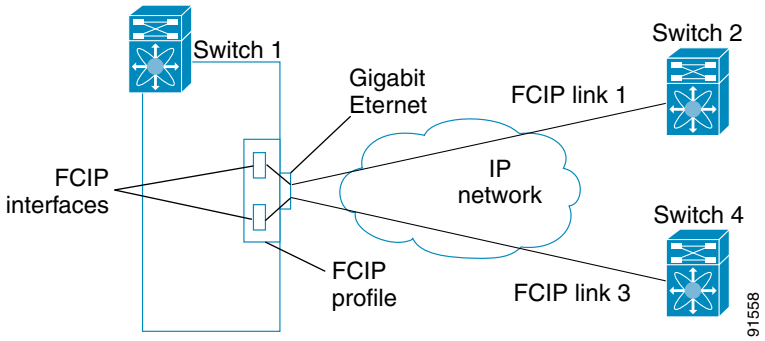

FCIP Interface

The FCIP interface is the local end point 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.

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 commands for the F IP feature are only available when FCIP is enabled on a switch. When you disable this feature, all related configurations are automatically discarded.

Basic FCIP Configuration

To configure a FCIP link, follow these steps on both switches.

Step 1 Configure the Gigabit Ethernet interface.

Step 2 Create a FCIP profile, assign the Gigabit Ethernet interface's IP address to the profile.

Step 3 Create a FCIP interface, 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

To create a FCIP profile, you must assign a local IP address of a Gigabit Ethernet interface or subinterface to the FCIP profile.

Figure 24-8 Assigning Profiles to Each Gigabit Ethernet Interface

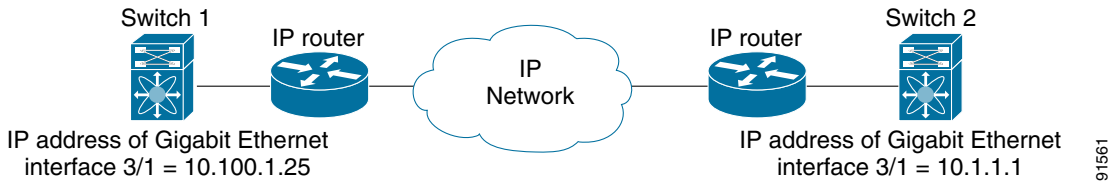

Creating FCIP Links

When two FCIP link end points are created, a FCIP link is established between the two IPS modules. To create a FCIP link, assign a profile to the FCIP interface and configure the peer information. The peer IP switch information initiates (creates) a FCIP link to that peer switch.

Figure 24-9 Assigning Profiles to Each Gigabit Ethernet Interface

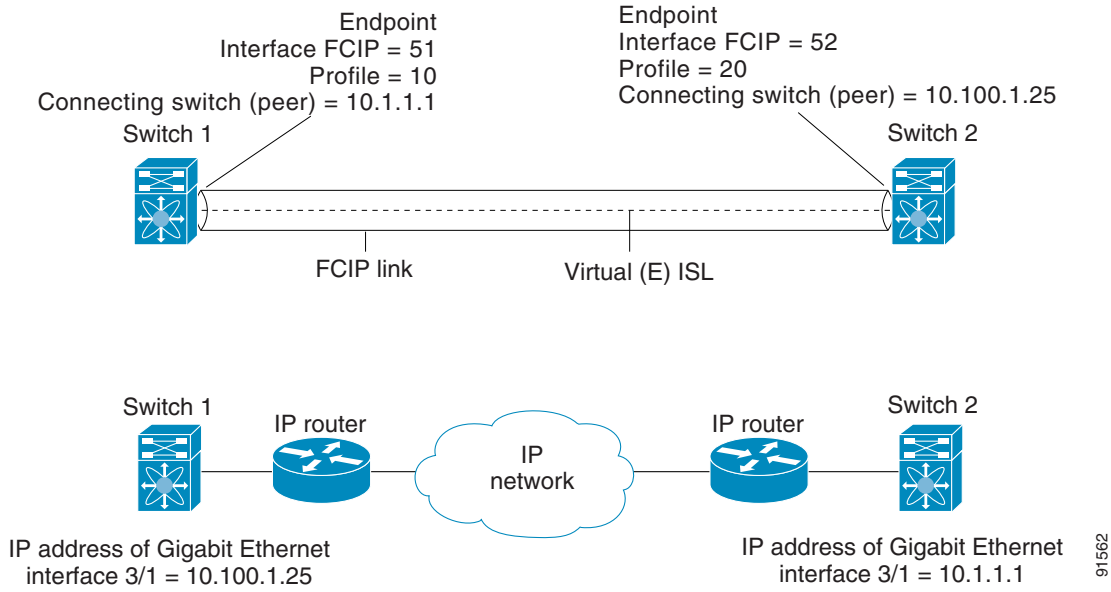

Creating FCIP Tunnels with Device Manager

To create and manage FCIP tunnels with Device Manager, first verify that the IPS module is inserted in the required Cisco MDS 9000 Family switches, and that the switches' Gigabit Ethernet interfaces are connected and the connectivity verified using the ping command. The steps in creating FCIP tunnels are:

- Assigning FCIP Profiles, page 24-11
- Creating Tunnels, page 24-11
- Verifying Interfaces, page 24-12

Assigning FCIP Profiles

You can use Device Manager to configure FCIP tunnels between switches. First, you must create FCIP profiles, and then bind the interfaces to the profile. To bind an FCIP profile to an interface, use the IP address of the interface in the FCIP profile's IP address configuration. Profile numbers range from 1 to 255. The interface associated with a profile can be either of the following:

- EtherChannel
- Ethernet subinterface slot and port (or slot, port, and VLAN ID)

To create and bind profiles on a Gigabit Ethernet interface, follow these steps.

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Open Device Manager.

Step 3 Choose **IP > FCIP**.

Step 4 Click the **Profiles** tab if it is not already selected. You see a list of any profiles that are already bound, along with their IP addresses, in the FCIP Profiles dialog box.

Step 5 To add a new profile, click **Create**. You see the Create FCIP Profiles dialog box.

Step 6 Enter the profile ID in the ID field.

Step 7 Select an IP address of the interface to which you want to bind the profile from the IP Address drop-down list.

Step 8 Enter all the optional information, if desired.

Step 9 Click **Create** to add this profile to the table. Click **Close** to close the Create FCIP Profiles dialog box without adding the profile.

Creating Tunnels

Each Gigabit Ethernet interface can have three active FCIP tunnels on it at one time.

To create these tunnels, follow these steps:

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Open Device Manager.

Step 3 Choose **IP > FCIP**.

Step 4 Click the **Tunnels** tab if it is not already selected. The FCIP Tunnels dialog box lists the remote IP address of the interface together with optional attributes.

Step 5 Click the **Create** button. You see the Create FCIP Tunnels dialog box.

Step 6 Enter the entity ID in the ID field.

Step 7 Enter a remote IP address as the endpoint to which you want to link.

Step 8 Enter all the optional information, if desired.

Step 9 Click **Create** to add this tunnel to the table. Click **Close** to close the Create FCIP Tunnels dialog box without adding the tunnel.

Verifying Interfaces

To verify the interfaces, follow these steps:

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Open Device Manager.

Step 3 Choose **Interface > FCIP**.

Step 4 Click the **Interfaces** tab (if it is not already selected) to see the FCIP Interfaces dialog box.

Verifying Extended Link Protocols (ELP)

To verify the extended link protocol, follow these steps:

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Open Device Manager.

Step 3 Choose **IP > FCIP**.

Step 4 Click the **ELP** tab (if it is not already selected) to see the FCIP ELP dialogist.

Checking Trunk Status

To check the trunk status, follow these steps:

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Open Device Manager.

Step 3 Choose **IP > FCIP**.

Step 4 Click the **Trunk Status** tab (if it is not already selected) to see the FCIP Trunk Status dialog box.

Checking for Interface Errors

To check for interface errors, follow these steps:

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Open Device Manager.

Step 3 Choose **IP > FCIP**.

Step 4 Click the **Interface Errors** tab (if it is not already selected) to see a list of FCIP-specific end-point/interface errors in the FCIP Interface Errors dialog box.

Creating FCIP Tunnels with Fabric Manager

To create and manage FCIP tunnels with Fabric Manager, you use the FCIP Wizard. First verify that the IPS module is inserted in the required Cisco MDS 9000 Family switches, and that the switches' Gigabit Ethernet interfaces are connected and the connectivity verified. The steps in creating FCIP tunnels using the FCIP Wizard are:

- Select the endpoints
- Choose the interfaces' IP addresses
- Specify link attributes

To create FCIP tunnels using the FCIP Wizard, follow these steps:

Step 1 Open the FCIP Tunnels Wizard by clicking on its icon in the Fabric Manager toolbar.

Figure 24-10 FCIP Tunnels Wizard

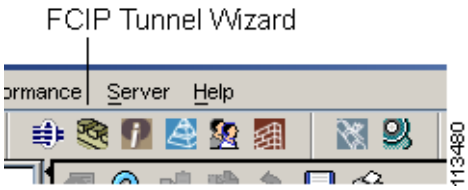

Step 2 Choose the endpoints.

Step 3 Choose the ports' IP addresses.

Step 4 Choose the link attributes.

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.

FCIP configuration options can be accessed from the switch(config- profile)# submode prompt.

Configuring TCP Listener Ports

The default TCP port for FCIP is 3225. You can change this port using the **port** command.

Configuring TCP Parameters

This section provides details on the TCP parameters that can be configured to control TCP behavior in a switch. The following TCP parameters can be configured.

Minimum Retransmit Timeout

The **tcp minimum-retransmit-time** option controls the minimum amount of time TCP waits before retransmitting. By default, this value is 200 milliseconds.

Keepalive Timeout

The **tcp keepalive-timeout** option enables you to configure the interval between which the TCP connection verifies if the FCIP link is functioning. This ensures that a 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. This command can be used to tune the time taken to detect FCIP link failures.

The first interval during which the connection is idle is 60 seconds (default). When the connection is idle for 60 seconds, 8 keepalive probes are sent at 1-second intervals. If no response is received for these 8 probes and the connection remains idle throughout, that FCIP link is automatically closed.

Only the first interval (during which the connection is idle) can be changed from the default of 60 seconds. This interval is identified using the **keepalive-timeout**

option. The valid range is from 1 to 7200 seconds.

Maximum Retransmissions

The **tcp max-retransmissions** option specifies the maximum number of times a packet is retransmitted before TCP decides to close the connection.

Path MTU

Path MTU (PMTU) is the minimum MTU on the IP network between the two end points 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 default timeout of 3600 seconds. If TCP reduces the size of the max segment because of PMTU change, the reset-timeout specifies the time after which TCP tries the original MTU.

SACK

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 time. 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 computed using the **max-bandwidth** option, the **min-available-bandwidth** option, and the dynamically-measured round-trip-time (RTT). The interaction and the resulting TCP behavior is outlined below:

The configured round-trip-time option determines the window scaling factor of the TCP connection. This option is only an approximation. The measured RTT value overrides the round-trip-time option 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.

- If the average rate of the fc traffic over the preceding RTT is less than the min-available-bandwidth * RTT, every FC burst is sent immediately at the min-available-bandwidth rate, provided no TCP drops occur.
- If the average rate of the FC traffic is greater than min-available-bandwidth * RTT, but less than max-bandwidth * RTT, then if the FC traffic is transmitted in burst sizes smaller than the configured CWM value all the bursts are sent immediately by FCIP at the max-bandwidth rate.
- If the average rate of the FC traffic is larger than the min-available-bandwidth * RTT and the burst size is greater than the CWM value, some traffic will not be sent immediately.

The maximum-bandwidth option and the measured round-trip-time together determine the maximum window size.

The **min-available-bandwidth** option and the measured **round-trip-time** together determine the threshold below which TCP aggressively maintains a window size sufficient to transmit at **min-available-bandwidth**. The software uses standard TCP rules to increase the window beyond the one required to maintain the **min-available-bandwidth** in order to reach the **max-bandwidth**. The defaults are **max-bandwidth** = 1G, **min-available-bandwidth** = 15 Mbps, and **round-trip-time** =1 ms.

Buffer Size

The **send-buffer-size** option defines 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 buffer size is 0 KB.

Quality of Service

The Quality of Service (QoS) feature specifies the differentiated services code point (DSCP) value to mark all IP packets (type of service—STOS 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.

Monitoring Window Congestion

The congestion window monitoring (CWM) option determines the maximum burst size allowed after an idle period.

- If the FC traffic burst is smaller than the configured CWM value, every packet is sent immediately, provided that no TCP drops were detected in the previous RTT.
- If FC traffic burst is larger than the configured CWM value, the excess packets will be sent during succeeding RTTs.

By default the **tcpcwm** option is enabled and the default burst size is 10 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.

Advanced FCIP Interface Configuration

You can establish connection to a peer by configuring one or more of the following options for the FCIP interface. To do so, you must first create the interface and enter the config-if submode.

Configuring Peers

To establish a FCIP link with the peer, you can use one of two options:

- —Used to configure both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.
- —Used to configure one end of the FCIP link when security gateways are present in the IP network. Optionally, you can also use the port 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.

Special Frames

You can alternatively establish a FCIP link with a peer using an optional protocol called special frames. You can enable or disable the **special-frame** option. On the peer side, the **special-frame** option must be enabled in order to establish the FCIP link. When the **special-frame** option is 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.

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.

Configuring Active Connection

Use the **passive-mode** option to configure the required mode for initiating an IP 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.

Ensure that both ends of the FCIP link are not configured as passive mode. If both ends are configured as passive, the connection will not be initiated.

Configuring the Number of TCP Connections

Use the **tcp-connection** option to specify the number of TCP connections from a FCIP link. By default, the switch tries two (2) TCP connections for each FCIP link. You can configure 1 or 2 TCP connections.

For example, the Cisco PA-FC-1G Fibre Channel port adapter which has only 1 (one) TCP connection interoperates with any switch in the Cisco MDS 9000 Family. One TCP connection is within the specified limit and you can change the configuration on the switch using the **tcp-connection 1** command. If the peer initiates one TCP connection, and your MDS switch is configured for two TCP connections, the software handles it gracefully and moves on with just one connection.

Enabling Time Stamps

Use the **time-stamp** option to enable or disable FCIP time stamps on a packet. The **time stamp** option instructs the switch to discard packets that are outside the specified time. By default, the **time-stamp** option is disabled.

The **acceptable-diff** option 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 if a packet arrives within a 1000 millisecond interval (+ or -1000 milliseconds), that packet is accepted.

If the **time-stamp** option is enabled, be sure to configure NTP on both switches.

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. depicts a typical SAN extension over an IP network.

Figure 24-11 FCIP B Port and Fibre Channel E Port

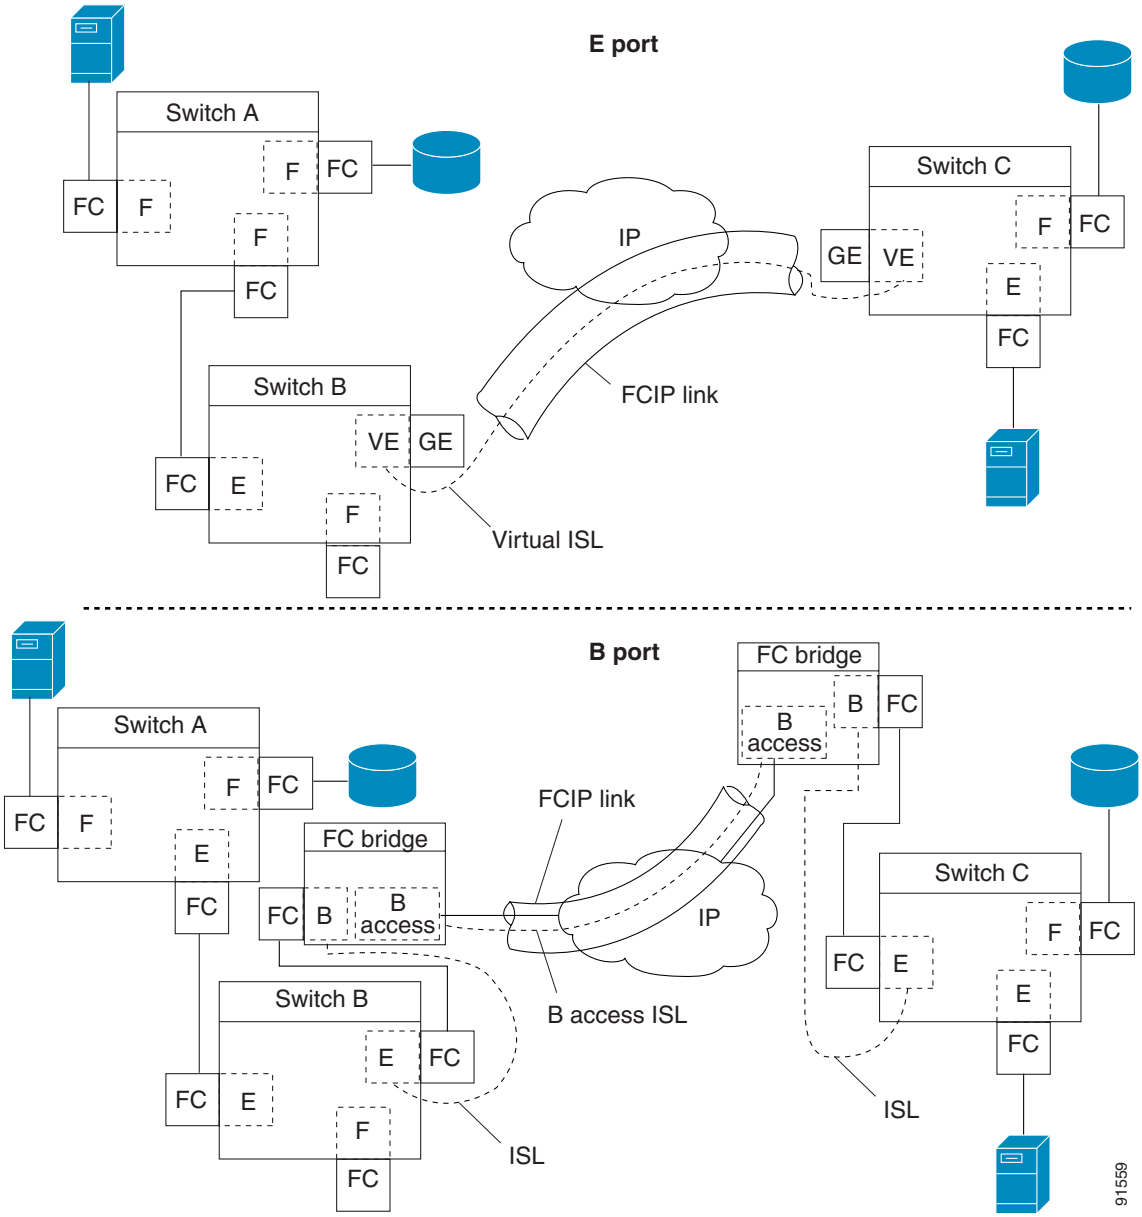

B ports bridge Fibre Channel traffic from one E port to a remote E port without participating in fabric-related activities such as principal switch election, Domain ID assignment, and Fibre Channel routing (FSPF). For example, Class F traffic entering a SAN extender does not interact with the B port. 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 which 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 a FCIP link, B ports use a B access ISL.

The IPS module supports 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 an virtual E port which completes the end-to-end E port connectivity requirement.

Figure 24-12 FCIP Link Terminating in a B Port Mode

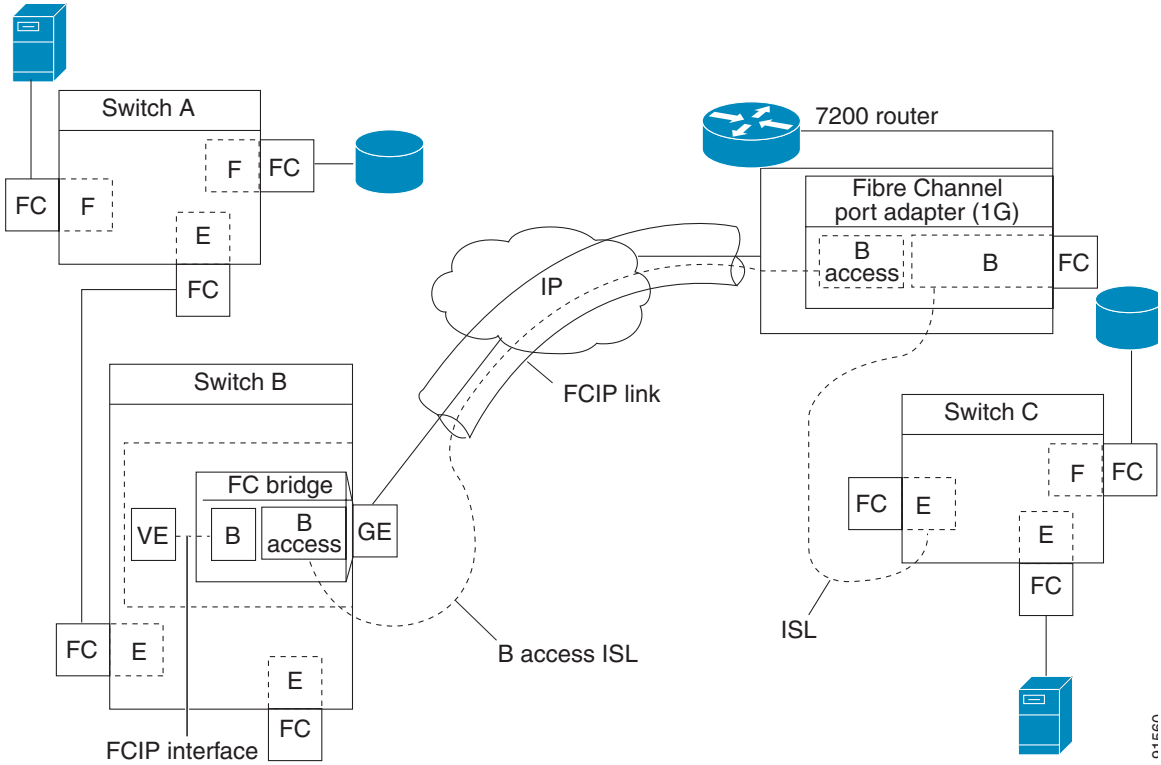

The B port feature in the IPS module allows remote B port SAN extenders to communicate directly with a Cisco MDS 9000 Family switch, therefore eliminating the need for local bridge devices.

Configuring B Ports

When a 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.

E Port Configurations

All configuration commands that apply to E ports, also apply to FCIP interfaces. The following features are also available FCIP interfaces:

- VSANs
  - FCIP interfaces can be a member of any VSAN.
  - Trunk mode
  - Trunk mode can be configured.
  - Trunk allowed VSANs can be configured

- 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
- Fibre Channel domains
- Zone merge
- Importing the zone database from the adjacent switch.
- Exporting the zone database from the adjacent switch.

Configuring FCIP Write Acceleration

The FCIP Write Acceleration feature in SAN-OS 1.3(3) enables you to significantly improve application 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 the command to transfer ready acknowledgement.

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 tunnel will not initialize.

Figure 24-13 FCIP Link Write Acceleration

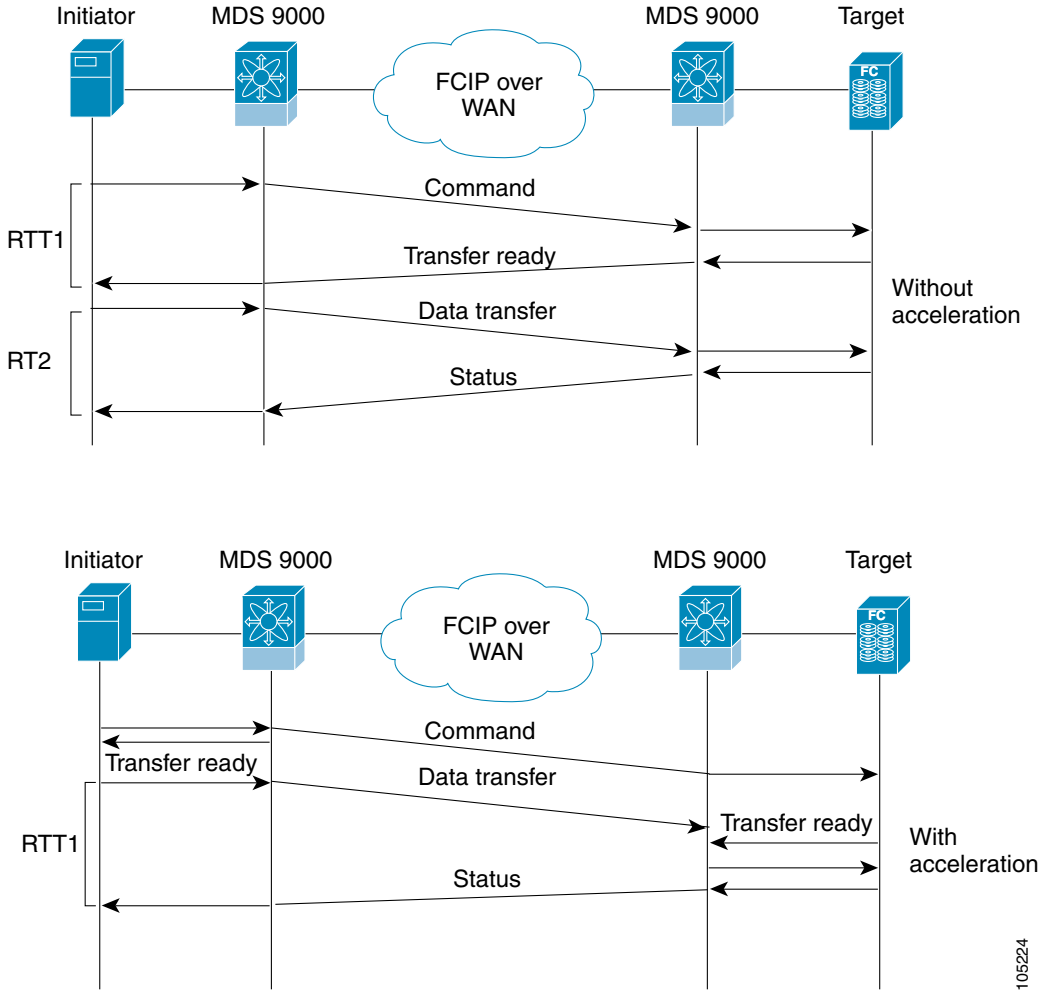

In Figure 24-13, some data sent by the host is queued on the target before the target issues a Transfer Ready. This way the actual write operation may be done in a less time than the write operation without the write acceleration feature being enabled.

Tip FCIP write acceleration will not work if the FCIP port is part of a PortChannel or if there are multiple paths with equal weight between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or broken write or read operations.

Enabling FCIP Compression

The FCIP compression feature introduced in Release 1.3(x) 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.

This feature uses the Lempel-Zif-Stac (LZS) compression algorithm to compress packets.

The *high-throughput* mode allows faster compression but the compression ratio may be lower. The *high-comp-ratio* mode allows a higher compression ratio, but the throughput may be lower.

Fibre Channel PortChannels

Figure 24-14 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 24-14 PortChannel Based Load Balancing

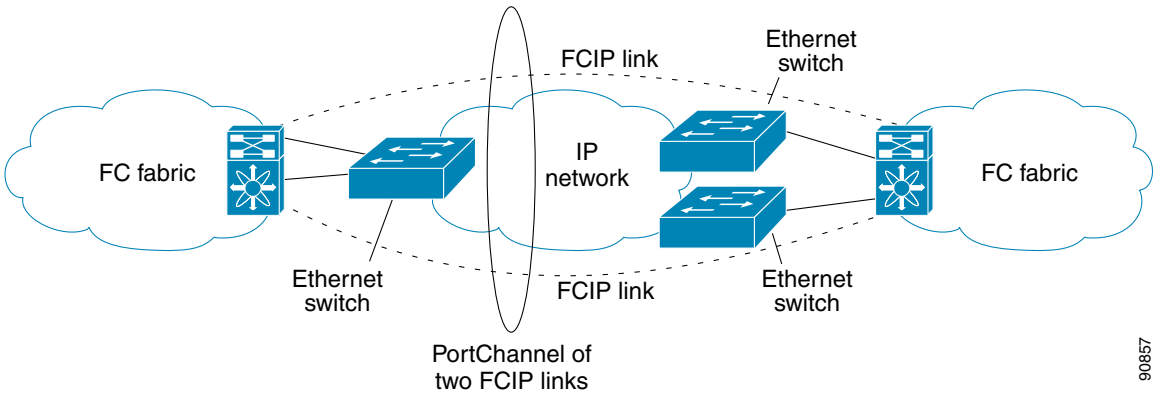

The following characteristics set Fibre Channel PortChannel solutions apart from other solutions:

- 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 24-15 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 24-15 FSPF-Based Load Balancing

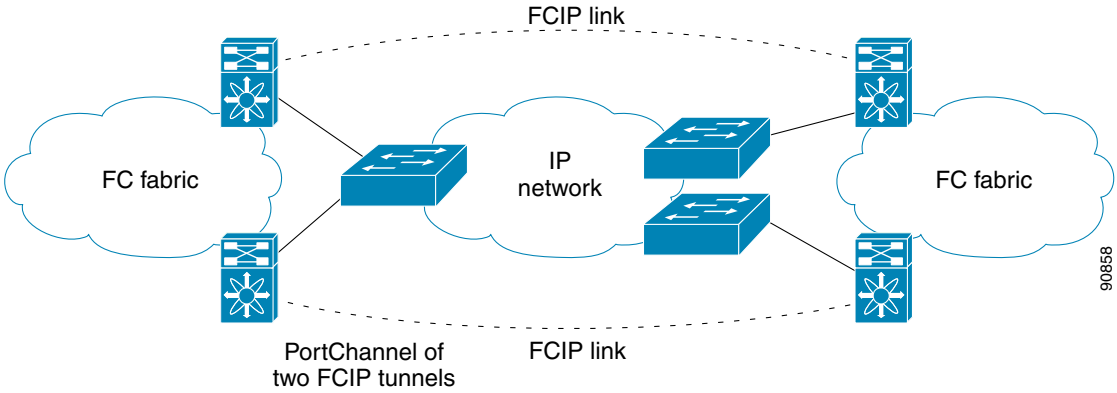

The following characteristics set FSPF solutions apart from other solutions:

- 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

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 24-16 VRRP-Based High Availability

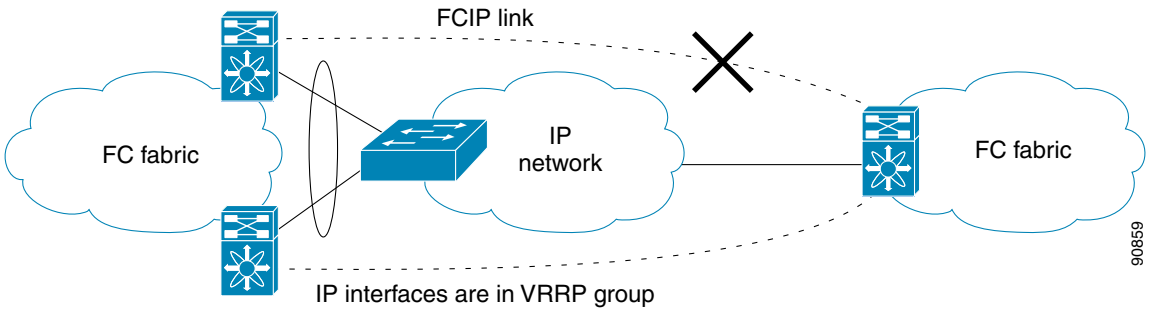

The following characteristics set VRRP solutions apart from other solutions:

- 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 24-17 displays a Ethernet PortChannel-based high availability FCIP example. This solution addresses the problem caused by individual Gigabit Ethernet link failures.

Figure 24-17 Ethernet PortChannel-Based High Availability

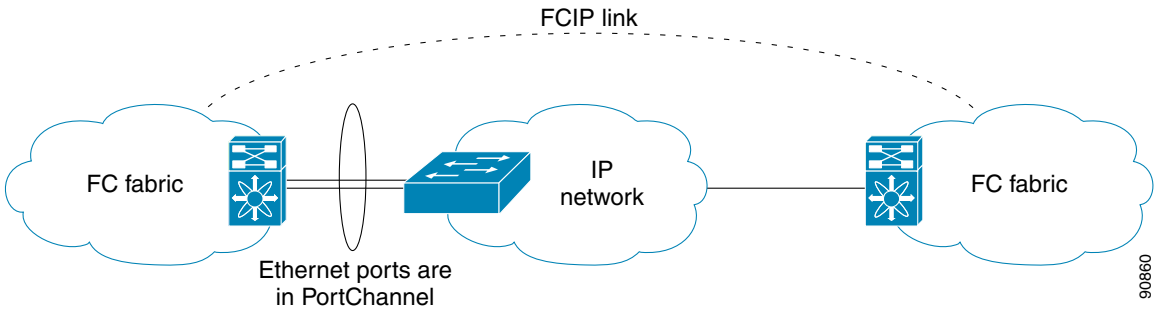

The following characteristics set Ethernet PortChannel solutions apart from other solutions:

- 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 appears like one logical Gigabit Ethernet link.
- The FCIP link stays up during the failover.

Ethernet PortChannels and Fibre Channel PortChannels

Ethernet PortChannels offer Ethernet-level redundancy, Fibre Channel PortChannels offer (E)ISL-level redundancy. FCIP is unaware of any Ethernet PortChannels or Fibre Channel PortChannels. Fibre Channel PortChannels are unaware of any Ethernet PortChannels, and there is no mapping between the two.

Figure 24-18 PortChannels at the Fibre Channel and Ethernet Levels

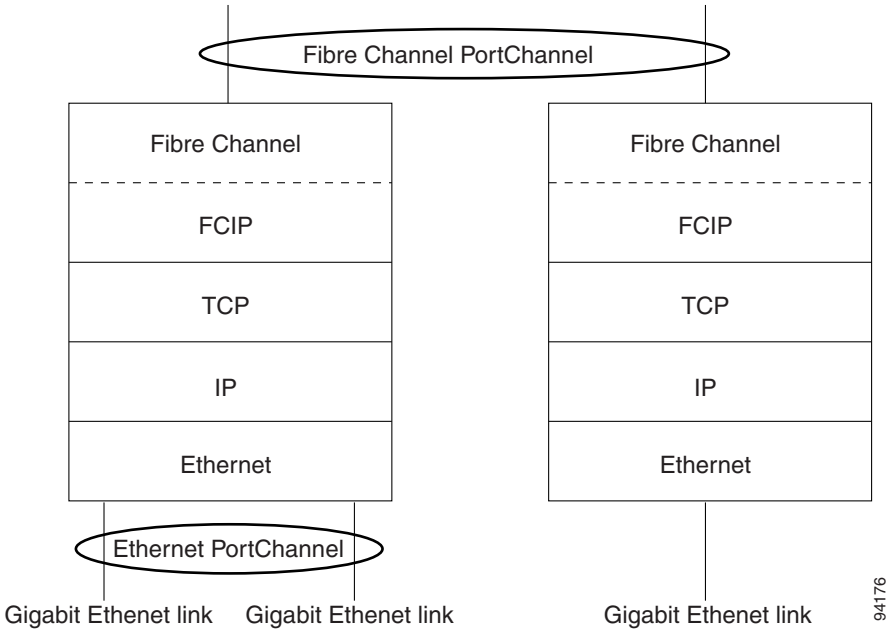

Configuring iSCSI

This section contains the following topics:

- About iSCSI, page 24-25
- Enabling iSCSI, page 24-26
- Routing iSCSI Requests and Responses, page 24-27
- Presenting Fibre Channel Targets as iSCSI Targets, page 24-27
- Presenting iSCSI Hosts as Virtual Fibre Channel Hosts, page 24-32
- Configuring iSCSI Proxy Initiators, page 24-34
- Access Control in iSCSI, page 24-36
- iSCSI User Authentication, page 24-37
- Advanced iSCSI Configuration, page 24-38
- iSCSI High Availability, page 24-39
- iSCSI Authentication Setup Guidelines, page 24-42

About iSCSI

The IPS module provides transparent SCSI routing by default. IP hosts using iSCSI protocol can transparently access targets on the Fibre Channel network. provides an example of a typical configuration of iSCSI hosts with access to a Fibre Channel SAN.

Figure 24-19 Typical IP to Fibre Channel SAN Configuration

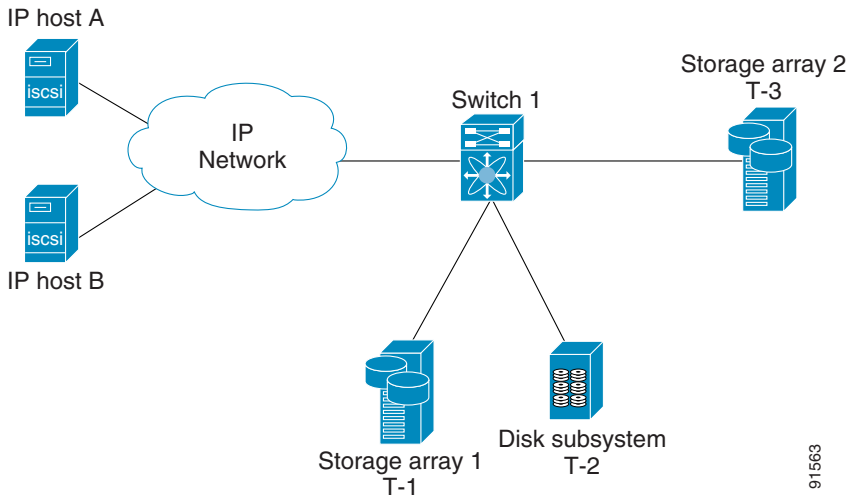

The IPS module enables you to create virtual iSCSI targets and maps them to physical Fibre Channel targets available in the Fibre Channel SAN. They present the Fibre Channel targets to IP hosts as if the physical targets were attached to the IP network.

Figure 24-20 iSCSI View

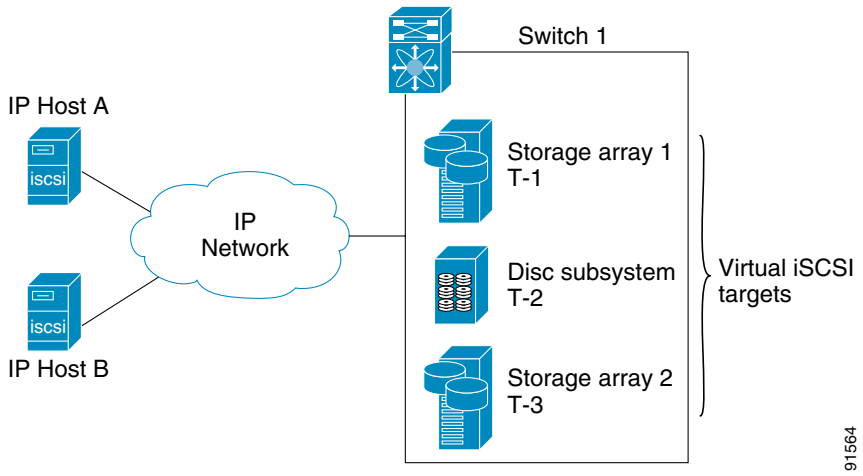

In conjunction with presenting Fibre Channel targets to iSCSI hosts, the IPS module presents each iSCSI host as a Fibre Channel host (in transparent mode), i.e. Host Bus Adaptor (HBA) to the Fibre Channel storage device. The storage device responds to each IP host as if it were a Fibre Channel host connected to the Fibre Channel network.

Figure 24-21 Fibre Channel SAN View

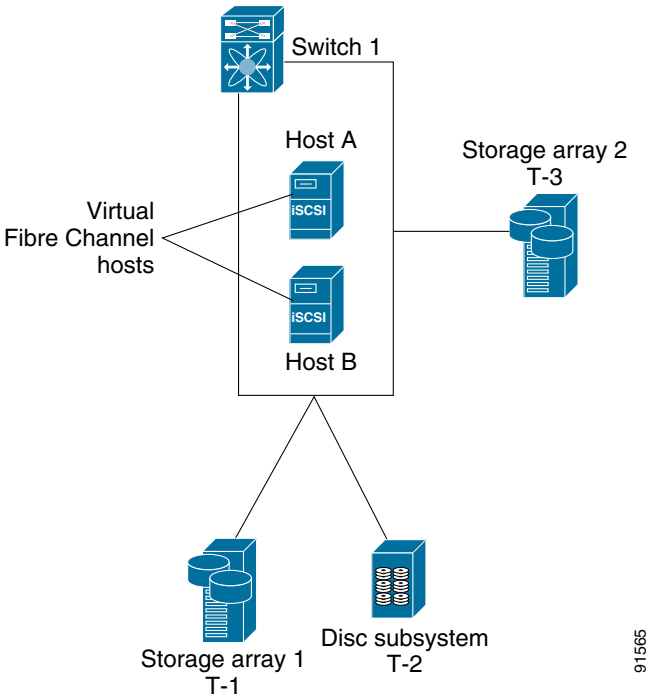

Refer to the IETF standards for IP storage at http://www.ietf.org, for information on the iSCSI protocol.

Enabling iSCSI

To begin configuring 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.

The configuration and verification commands for the iSCSI feature are only available when iSCSI is enabled on a switch. When you disable this feature, all related configurations are automatically discarded.

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

Step 1 Choose **End Devices > ISCSI** from the Physical Attributes pane. The ISCSI tables display in the Information pane.

Step 2 Click the **Control** tab if it is not already displayed. This shows the iSCSI enable status for all switches in the fabric that contain IPS ports

Step 3 Choose Enable from the Command column for each switch that you want to enable iSCSI on.

Step 4 Click the **Apply Changes** icon to save these changes or click the **Undo Changes** icon to remove all changes without saving them.

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 Choose an existing iSCSI initiator or add the iSCSI node name or IP address for a new iSCSI initiator.

Step 3 Choose the switch for this iSCSI initiator if you are adding a new iSCSI initiator and click **Next**.

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

Step 5 Set the zone name for this new iSCSI zone and optionally check the **Read Only** check box.

Step 6 Click **Finish** to create this iSCSI initiator, or click **Cancel** to close the wizard without creating the iSCSI initiator.

Routing iSCSI Requests and Responses

The iSCSI feature consists of routing iSCSI requests and responses between 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.

Each iSCSI host that requires access to storage via the IPS module needs to have a compatible iSCSI driver installed. (The CCO website at provides a list of compatible drivers). Using 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 for a peripheral channel in the host. From the storage device perspective, each IP host appears as a Fibre Channel host.

Routing SCSI from the IP host to the Fibre Channel storage device consists of the following main actions.

- Transporting iSCSI requests and responses over an IP network between hosts and the IPS module.
- Routing SCSI requests and responses between hosts on an IP network and the Fibre Channel storage device (converting iSCSI to FCP and vice versa). This routing is performed by the IPS module.
- Transporting FCP requests or responses between the IPS module and Fibre Channel storage devices.

FCP (the Fibre Channel equivalent of iSCSI) carries SCSI commands over a Fibre Channel SAN.

Presenting Fibre Channel Targets as iSCSI Targets

The IPS module presents physical Fibre Channel targets as iSCSI targets allowing them to be accessed by iSCSI hosts. It does this in one of two ways:

- —Used if all logical units (LUs) in all Fibre Channel storage targets are made available to iSCSI hosts (subject to VSAN and zoning).
- —Used if iSCSI hosts are restricted to subsets of LUs in the Fibre Channel targets and additional iSCSI access control is needed. Also, static import allows automatic failover if the Fibre Channel target LU is reached by redundant Fibre Channel ports.

The IPS module does not import Fibre Channel targets to iSCSI by default. Either dynamic or static mapping must be configured before the IPS module makes Fibre Channel targets available to iSCSI initiators. When both are configured, statically mapped Fibre Channel targets have a configured name. Targets that are not statically imported are advertised with the name created by the conventions explained in this section.

Dynamic Importing

To enable dynamic importing of Fibre Channel targets into iSCSI, use the **iscsi import target fc** command.

The IPS module maps each physical Fibre Channel target port as one iSCSI target. That is, all LU accessible via the physical storage target port are available as iSCSI LUs with the same LU number (LUN) as in the storage target.

For example, if an iSCSI target was created for Fibre Channel target port with pWWN 31:00:11:22:33:44:55:66 and that pWWN contains LUN 0 through 2, those LUNs would become available to an IP host as LUNs 0 through 2 as well.

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

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 creates an IQN formatted iSCSI node name using the following conventions:

- IPS ports that are not part of a VRRP group 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:05.<mgmt-ip-address>.pc-<port-ch-sub-intf#>.<Target-pWWN>
  ```

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

Configuring Dynamic Importing with Device Manager

To dynamically import Fibre Channel targets as iSCSI targets, follow these steps:

Step 1 Choose **IP > iSCSI** from Device Manager. You see the iSCSI dialog box.

Step 2 Choose the **Targets** tab to display a list of existing iSCSI targets.

Step 3 Check the **Dynamically Import FC Targets** check box.

Step 4 Click **Apply** to save this change or click **Cancel** to close the dialog box without saving any changes.

Static Importing

You can manually (statically) create an iSCSI target and assign a node name to it. A statically-mapped iSCSI target can either contain the whole FC target port, or it can contain one or more LUs from a Fibre Channel target port.

To create a static iSCSI virtual target for the Fibre Channel target port, follow these steps

Step 1 Choose **IP > iSCSI** from Device Manager. You see the iSCSI dialog box.

Step 2 Choose the **Targets** tab to display a list of existing iSCSI targets.

Step 3 Click **Create** to create a new 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 for the Fibre Channel target port you are mapping.

Step 6 Click the **List** radio button and set the iSCSI initiator node names or IP Addresses that you want this virtual iSCSI target to access, or click the **All** radio button to let the iSCSI target access all iSCSI initiators. See the "Access Control in iSCSI" section on page 24-36.

Step 7 Click 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.

For multiple interfaces configured with iSNS, a different static virtual target name has to be created for each interface tagged to an iSNS profile and each static virtual target must be advertised only from one interface.

Advertising iSCSI Targets

You can limit the Gigabit Ethernet interfaces over which static iSCSI targets are advertised. By default iSCSI targets are advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces.

High Availability Static Target Importing

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 24-22).

Figure 24-22 Static Target Importing Through Two Fibre Channel Ports

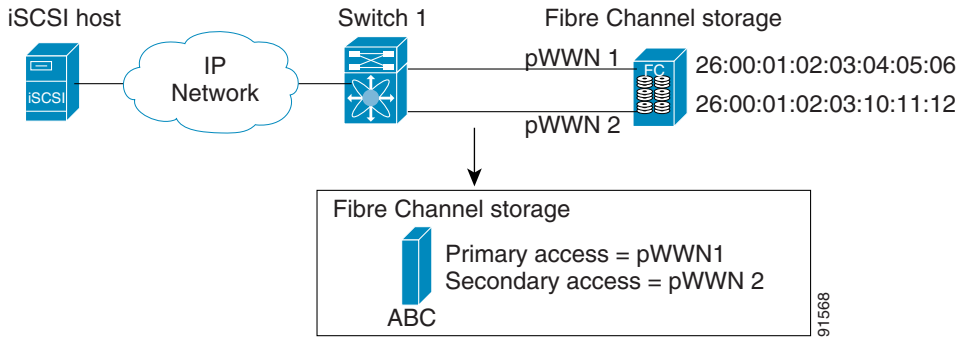

In Figure 24-22, you can create a virtual iSCSI target that is mapped to both pWWN1 and pWWN2 to provide redundant access to the Fibre Channel targets.

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

Refer to the *Cisco MDS 9000 Family Configuration Guide* for details on setting the secondary pWWN.

Enable the **revert to 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 does 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 enable the **revert to primary port** option, follow these steps:

Step 1 Choose **End Devices > iSCSI** from the Physical Attributes pane in Fabric Manager, or choose **IP > iSCSI** from Device Manager. You see the iSCSI tables in the Fabric Manager Information pane, or the Device Manager iSCSI dialog box.

Step 2 Click the **Targets** tab to display a list of existing iSCSI targets.

Step 3 Check the **RevertToPrimaryPort** check box to enable this option.

Step 4 Set the iSCSI target node name in the iSCSI Name field, in IQN format.

Step 5 Click the **Apply Changes** icon in Fabric Manager (or click **Apply** in Device Manager) to save this change, or click **Cancel** to close the dialog box without saving any changes.

Configuring the Trespass Feature

In addition to the high availability of statically imported iSCSI targets, the trespass feature is available (effective Release 1.3(x)) to enable the export 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** command be issued, to export 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 MDS issues a trespass command to the target to export the LUs on the new active port. The iSCSI session switches to use the new active port and the exported LUs are accessed over the new active port.

Figure 24-23 Virtual Target with an Active Primary Port

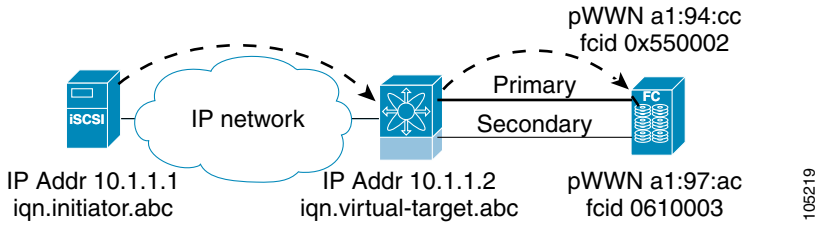

To configure the trespass feature, follow these steps:

Step 1 Choose **End Devices > iSCSI** from the Physical Attributes pane in Fabric Manager, or choose **IP > iSCSI** from Device Manager. You see the iSCSI tables in the Fabric Manager Information pane, or the Device Manager iSCSI dialog box.

Step 2 Click the **Targets** tab to display a list of existing iSCSI targets.

Step 3 Check the **Trespass Mode** check box to enable this option.

Step 4 Click the **Apply Changes** icon in Fabric Manager (or click **Apply** in Device Manager) to save this change, or click **Cancel** to close the dialog box without saving any changes.

iSCSI Virtual Target Configuration Examples

This section provides three examples of virtual target configurations.

Example 1

This example assigns the whole Fibre Channel target as a virtual iSCSI target. All LUNs that are part of the Fibre Channel target are available as part of the iSCSI target.

Figure 24-24 Assigning iSCSI Node Names

Example 2

This example maps a subset of LUNs of a Fibre Channel target to three iSCSI virtual targets. Each iSCSI target only has one LUN.

Figure 24-25 Mapping LUNs to iSCSI a Node Name

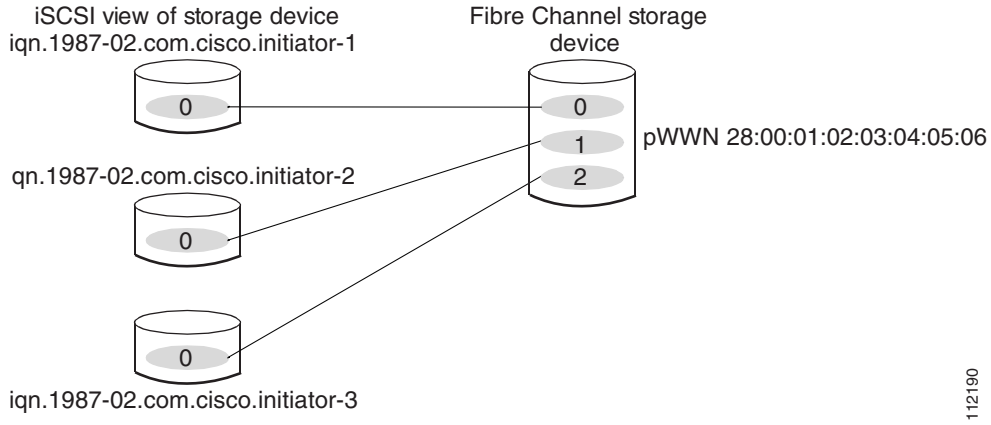

Example 3

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.

Figure 24-26 Mapping LUNs to Multiple iSCSI Node Names

Presenting iSCSI Hosts as Virtual Fibre Channel Hosts

The iSCSI hosts are mapped to virtual Fibre Channel hosts in one of two ways:

- (default)—Used if no access control is done on the Fibre Channel target. An iSCSI host may use different pWWNs each time it connects to a Fibre Channel target.
- —Used if an iSCSI host should always have the same pWWN or nWWN each time it connects to a Fibre Channel target.

Dynamic Mapping

When an iSCSI host connects to the IPS module using the iSCSI protocol, a virtual N port is created for the host. The nWWNs and pWWNs are dynamically allocated from the switch Fibre Channel WWN pool. The IPS module registers this N port in the Fibre Channel SAN. The IPS module continues using that nWWN and pWWN to represent this iSCSI host until it no longer has a connection to any iSCSI target via that IP storage port.

At that point, the virtual Fibre Channel host is taken offline from the Fibre Channel SAN and the nWWNs and pWWNs are released back to the switch's Fibre Channel WWN pool. These addresses becomes available for assignment to other iSCSI hosts requiring access to Fibre Channel SANs.

When a dynamically mapped iSCSI initiator has multiple sessions to multiple Fibre Channel targets, each session can use the same pWWN and nWWN as long as it uses the same node name in the iSCSI login message.

Identifying Initiators

An iSCSI initiator is identified in one of two ways:

- By iSCSI node name—An initiator with multiple IP addresses (multiple interface cards--NICs or multiple network interfaces) has one virtual N port, assuming it uses the same iSCSI initiator name to iSCSI targets from all interfaces.
- By IP address—A virtual N port is created for each IP address it uses to login to iSCSI targets.

By default, the switch uses the iSCSI node name to identify the initiator.

Static Mapping

With dynamic mapping, each time the iSCSI host connects to the IPS module a new Fibre Channel N port is created and the nWWNs and pWWNs allocated for this N port may be different. Use the static mapping method to obtain the same nWWN and pWWNs for the iSCSI host each time it connects to the IPS module.

You can implement static mapping in one of two ways: system assignment or manual assignment.

- System assignment—When a static mapping configuration is created, one nWWN and/or one or more pWWNs are allocated from the switch's Fibre Channel WWN pool and the mapping is kept permanent.
- Manual assignment—You can specify your own unique WWN by providing them during the configuration process.

Tip We recommend using the **system assignment** option. If you manually assign a WWN, you must ensure its uniqueness.

Static mapping can be used on the IPS module to access intelligent Fibre Channel storage arrays that have access control and LUN mapping/masking configuration based on the initiator's pWWNs and/or nWWNs.

If an iSCSI host connects to multiple IPS ports, each port independently creates one virtual N port for the host. If static mapping is used, enough pWWNs should be configured for as many IPS ports to which a host connects.

If a system assignment option is used to configure WWNs for an iSCSI initiator, when the configuration is backed up to an ASCII file the system-assigned WWNs are also saved. Subsequently if you issue a CLI **write erase** command, you must manually delete the WWN configuration from the ASCII file.

Making the Dynamic Initiator WWN Mapping Static

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

Assigning VSAN Membership to iSCSI Hosts

By default, a host is only in VSAN 1 (default VSAN). You can configure an iSCSI host to be a member of one or more VSANs. The IPS module creates one Fibre Channel virtual N port in each VSAN to which the host belongs.

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.

Assigning VSANs to a iSCSI Interface

All dynamic iSCSI initiators are members of VSAN 1. 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, but can be changed. All dynamic iSCSI initiators are member of the port VSAN of the iSCSI interface.

Tip This is a 1.3(x) feature. If you downgrade to an earlier release, be sure to delete any assigned VSAN before performing the downgrade procedure.

To modify the VSANs assigned to an iSCSI interface using Device Manager, follow these steps:

Step 1 Choose **Interfaces > Ethernet or iSCSI**. You see the interfaces dialog box.

Step 2 Click the **iSCSI** tab. You see the iSCSI interface configuration table.

Step 3 Double-click on the PortVSAN column to modify the default port VSAN.

Step 4 Click **Apply** to save these changes, or click **Cancel** to discard changes.

Configuring iSCSI Proxy Initiators

When an interface is in the proxy initiator mode, you can only configure Fibre Channel access control (zoning) based on the Fibre Channel interface attributes--the WWN pair and available FCIDs. You cannot configure zoning using iSCSI attributes such as the IP address or the iQN name of the iSCSI initiator. To enforce initiator-based access control, use iSCSI based access control.

By default, each iSCSI initiator appears as one Fibre Channel initiator in transparent mode in the Fibre Channel fabric. For some storage arrays, this appearance requires the initiator's pWWN to be manually configured for access control purposes. This process can be quite cumbersome. The Proxy initiator feature allows all iSCSI initiators to connect through one IPS port making it appear as one Fibre Channel

port per VSAN. It simplifies the task of configuring the pWWN for each new initiator on the storage array, and Fibre Channel access control such as zoning. This feature along with static target importing (using LUN mapping) results in the configuration being performed only on the switch when a new iSCSI host is added. On the storage array, all LUNs that will be used by iSCSI initiators are configured to allow access by the proxy initiator's pWWN. From the iSCSI perspective, this configuration is no different from the default mode.

Figure 24-27 The iSCSI View of a Proxy Initiator

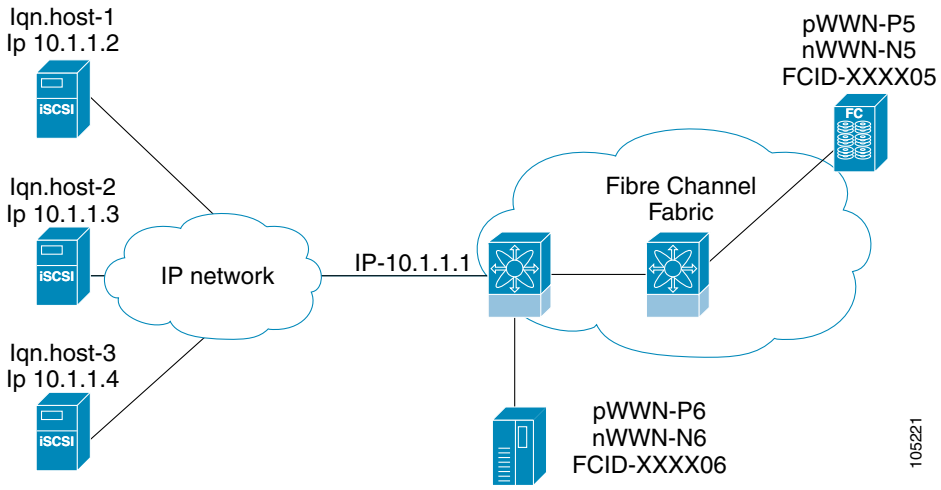

From the Fibre Channel perspective, only one Fibre Channel initiator is visible per VSAN.

Figure 24-28 The FC View with a Proxy Initiator

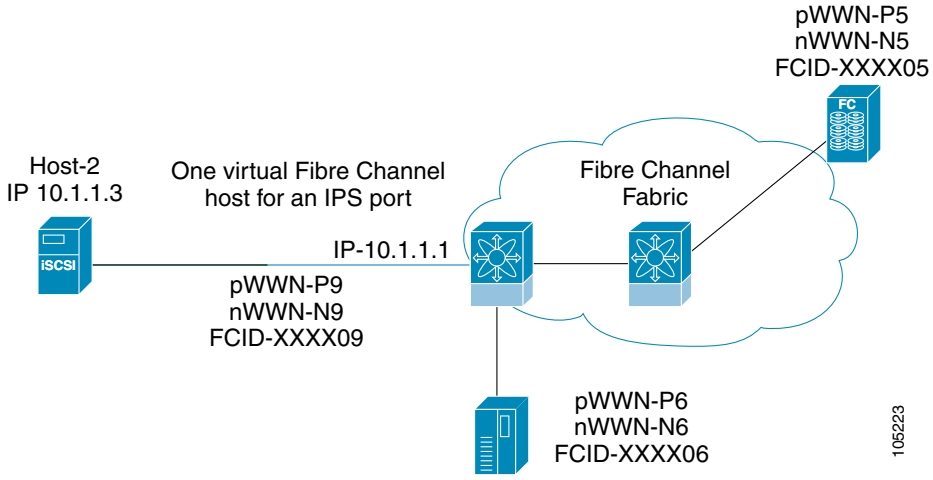

Configuring the iSCSI Proxy Initiator

To configure the proxy initiator, follow these steps:

Step 1 Choose Interfaces **> FC Logical** from the Fabric Manager Physical Attributes pane, or choose **Interfaces > Ethernet or iSCSI** from Device Manager. You see the Interface tables in the Fabric Manager Information pane or the Device Manager Interfaces dialog box.

Step 2 Click the **iSCSI** tab. You see the iSCSI interface configuration table.

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.

Access Control in iSCSI

You can control access to each statically-mapped iSCSI target by specifying a list of IPS ports on which it will be advertised and specifying a list of iSCSI initiator node names allowed to access it. Fibre Channel zoning-based access control and iSCSI-based access control are the two mechanisms by which access control can be provided for iSCSI. Both methods can be used simultaneously.

This access control is in addition to the existing Fibre Channel access control. The iSCSI initiator has to be in the same VSAN and zone as the physical Fibre Channel target.

Fibre Channel Zoning-Based Access Control

Zoning is an access control mechanism within a VSAN. The switch zoning implementation extends the VSAN and zoning concepts from the Fibre Channel domain to also cover the iSCSI domain. This extension includes both iSCSI and Fibre Channel features and provides a uniform, flexible access control across a SAN. Static and dynamic are the two Fibre Channel zoning access control mechanisms.

- Static—Statically map the iSCSI host to Fibre Channel virtual N port(s). This creates a permanent nWWNs and pWWNs. Next, configure the assigned pWWN into zones, similar to adding a regular Fibre Channel host pWWN to a zone.
- Dynamic—Add the iSCSI host's initiator node name as a member of a zone. When the IP host's Fibre Channel virtual N port is created and the Fibre Channel address (nWWNs and pWWNs) is assigned, Fibre Channel zoning is enforced.

To register an iSCSI host initiator as a member of a zone using Fabric Manager, follow these steps:

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

Step 2 Choose the VSAN and zone to which you want to add the iSCSI host initiator.

Step 3 From the list of available devices including iSCSI host initiators, click on the initiators that you want to add to the zone and click **Add to Zone or Alias**.

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

iSCSI-Based Access Control

For static iSCSI targets, you can manually configure a list of iSCSI initiators that are allowed to access it. The iSCSI initiator is identified by the iSCSI node name or the IP address of the iSCSI host.

By default, static virtual iSCSI targets are not accessible to any iSCSI host. You must explicitly configure accessibility to allow a virtual iSCSI target to be accessed by all hosts. The initiator access list can contain one or more initiators. Each initiator is identified by one of the following:

- iSCSI node names
- IP addresses
- IP subnets

Enforcing Access Control

IPS modules use both iSCSI node name-based and Fibre Channel zoning-based access control lists to enforce access control during iSCSI discovery and iSCSI session creation.

- iSCSI discovery--When an iSCSI host creates an iSCSI discovery session and queries for all iSCSI targets, the IPS 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.
- iSCSI session creation--When an IP host initiates an iSCSI session, the IPS module verifies if the specified iSCSI target (in the session login request) is a static mapped target, and if true, verifies if the IP host's iSCSI node name is allowed to access the target. If the IP host does not have access, its login is rejected.

The IPS module, then creates a Fibre Channel virtual N port (the N port may already exist) for this IP host and does a Fibre Channel name server query for the FCID of the Fibre Channel target pWWN that is being accessed by the IP host. It uses the IP host virtual N port's pWWN as the requester of the name server query. Thus, the name server does a zone-enforced query for the pWWN and responds to the query.

If the FCID is returned by the name server, then the iSCSI session is accepted. Otherwise, the login request is rejected.

iSCSI User Authentication

The IPS module supports the iSCSI authentication mechanism to authenticate iSCSI hosts that request access to storage. When iSCSI authentication is enabled, the iSCSI hosts must provide user name and password information each time an iSCSI session is established.

Only the Challenge Handshake Authentication Protocol (CHAP) authentication method is supported.If no authentication is configured, local authentication is used. You can use RADIUS authentication or TACACS+ authentication.

Configuring an Authentication Mechanism

During an iSCSI login, both the iSCSI initiator and target have the option to authenticate each other. By default, the IPS module allows either CHAP authentication or no authentication from iSCSI hosts.

Note The authentication for a Gigabit Ethernet interface or subinterface configuration overrides the authentication for the global interface configuration.

To configure an authentication method for iSCSI, follow these steps:

Step 1 Choose **End Devices > iSCSI** from the Fabric Manager Physical Attributes pane, or choose **IP > iSCSI** from Device Manager. You see the iSCSI tables in the Fabric Manager Information pane, or the Device Manager iSCSI dialog box.

Step 2 Click the **Global** tab. You see the iSCSI authentication configuration table.

Step 3 From Fabric Manager, choose chap or none from the authMethod column.

From Device manager, check the **Chap** check box to configure DH-CHAP authentication, or check the **none** check box for no authentication.

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.

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.

Advanced iSCSI Configuration

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

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

- iSCSI listener port—Configure the TCP port number for the iSCSI interface which listens for new TCP connections. The default port number is 3260. Following that, the iSCSI port only accepts TCP connections on the newly configured port
- TCP tuning parameters—The following TCP parameters can be configured.
  - SACK is enabled by default for iSCSI TCP configurations.
  - QoS configurations differ for iSCSI and FCIP interfaces.
- Identification of dynamic iSCSI initiator—iSCSI initiators are identified based on their IQN name or their IP address. In the absence of any configuration for the initiator (WWN or VSAN membership), the identifier key is the default connection. By default, the key is the IQN name but can be changed to IP address by toggling this mode.
- Proxy or transparent Initiator—For each iSCSI initiator with iSCSI target sessions, the switch creates a virtual FC initiator with a distinct pair of WWNs per VSAN. Targets that have access control per LUN, the WWN pair of each FC initiator must be configured in the target. The proxy initiator mode can be enabled to facilitate this configuration, in which case, all iSCSI initiators which connect to this iSCSI interface inherit the same WWN pair and create only one virtual FC initiator in each VSAN.

Setting the QOS Values

To set the QoS values, follow these steps:

Step 1 Choose **Interfaces > FC Logical** from the Fabric Manager Physical Attributes pane, or choose **Interfaces > Ethernet or iSCSI** from Device Manager. You see the Interface tables in the Fabric Manager Information pane, or the Device Manager Interfaces dialog box.

Step 2 Click the **iSCSI TCP** tab. 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 Forwarding Mode

The iSCSI gateway on the IPS module has two modes of forwarding operation

- The *pass-thru* mode (default)—In this mode, the IPS port converts an iSCSI PDU into an FCP frame or vice versa and then forwards it one frame or PDU at a time. The absence of buffering PDUs or frames keeps the operation latency low. To operate in this mode, the IPS port has to negotiate with its peers a suitable maximum size of the data payload in each frame/PDU. This is done during iSCSI login and FC PLOGI and the value is restricted by the TCP connection's Maximum Segment Size (MSS) and the maximum Fibre Channel data payload size specified by the FC target. This usually results in a smaller maximum payload size than most hosts expect, thus comes the second mode of forwarding.
- The *store-and-forward* mode—This mode allows the iSCSI client to send and receive an iSCSI data payload at the size it desires. This sometimes results in better performance for the client. The IPS port stores each TCP segment it receives until one full iSCSI PDU is received before converting and forwarding it as Fibre Channel frames to the FC target. In the opposite direction, the IPS port assembles all FC data frames of an exchange to build one iSCSI data-in PDU before forwarding it to the iSCSI client. The limitation on this mode is iSCSI CRC data digest cannot be used.

iSCSI High Availability

The following high availability features are available for iSCSI configurations:

- Multiple IPS Ports Connected to the Same IP Network
- VRRP-Based High Availability
- Ethernet PortChannel-Based High Availability

Multiple IPS Ports Connected to the Same IP Network

Figure 24-29 provides an example of a configuration with multiple Gigabit Ethernet interfaces in the same IP network.

Figure 24-29 Multiple Gigabit Ethernet Interfaces in the Same IP Network

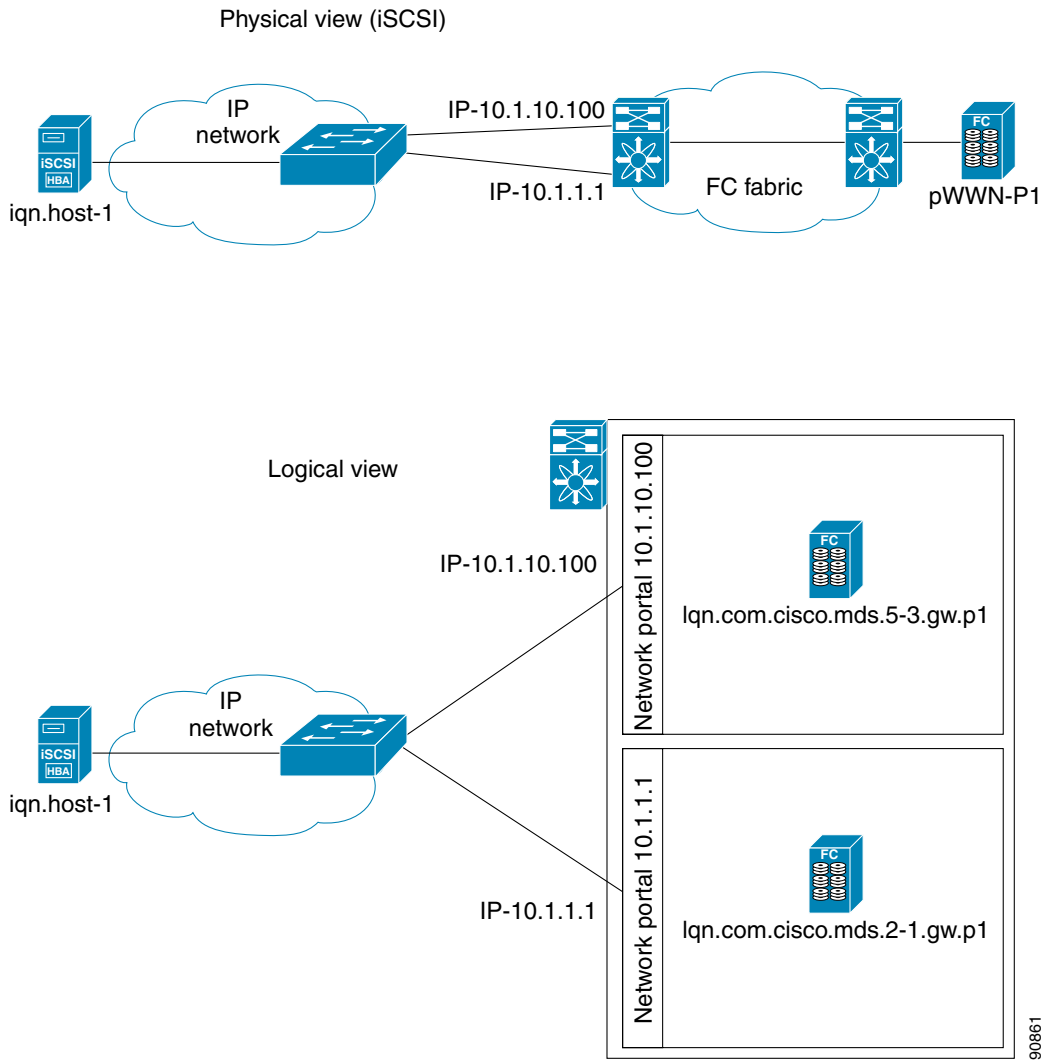

In Figure 24-29, each iSCSI host discovers two iSCSI targets for every physical Fibre Channel target (with different names). The multi-pathing software on the host provides load-balancing over both paths. If one Gigabit Ethernet interface fails, the host multi-pathing software is not affected because it can use the second path.

VRRP-Based High Availability

Figure 24-30 provides an example of a VRRP-based high availability iSCSI configuration.

Figure 24-30 VRRP-Based iSCSI High Availability

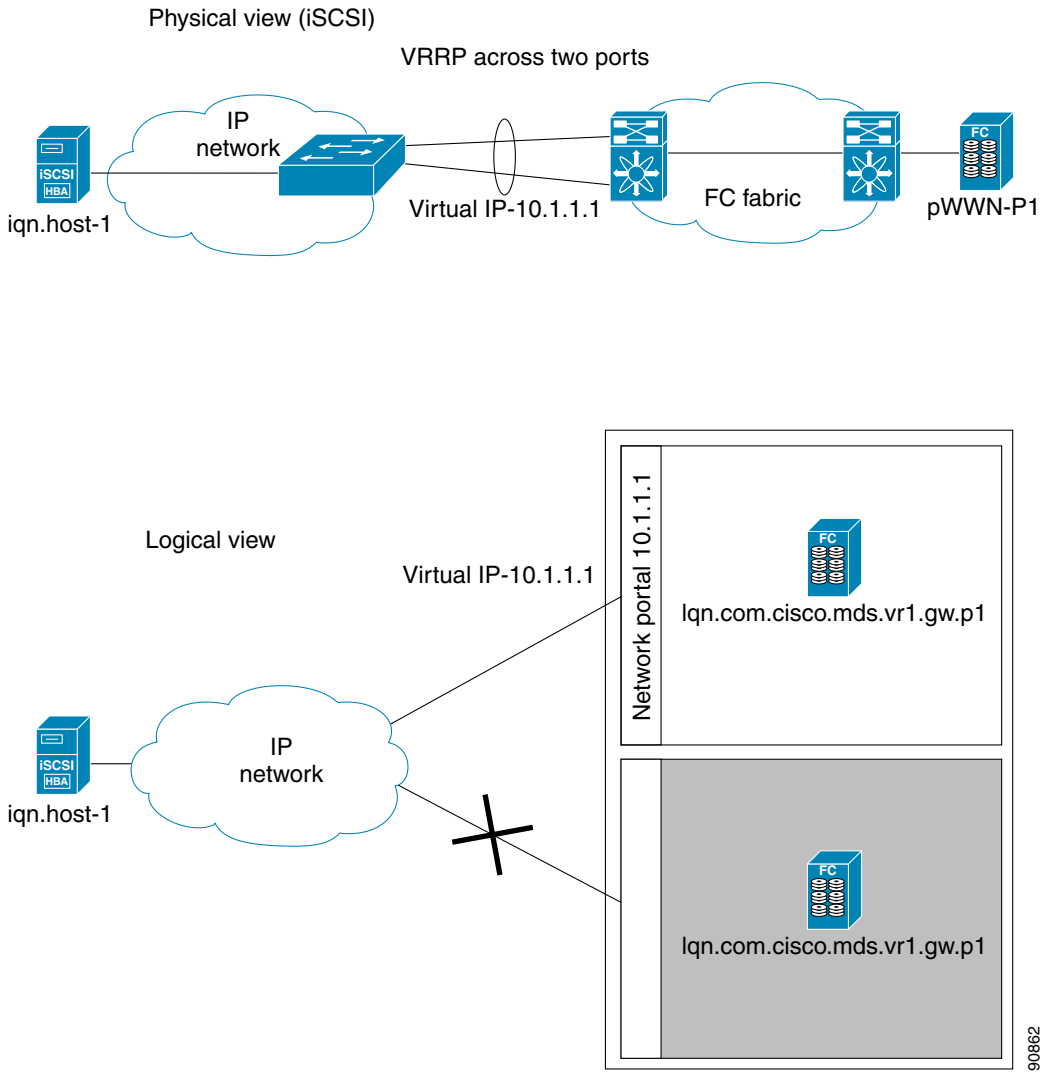

In Figure 24-30, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. When the Gigabit Ethernet interface of the VRRP master fails, the iSCSI session is terminated. The host then reconnects to the target and the session comes up because the second Gigabit Ethernet interface has taken over the virtual IP address as the new master.

Tip Ports that act as VRRP master and backup can be on different switches. If you have a static WWN configuration for iSCSI initiators, configure a different WWN for the iSCSI initiator for each switch. If you use a proxy-initiator, be sure to configure a different pWWN on each iSCSI interface for each VRRP port used.

Ethernet PortChannel-Based High Availability

All iSCSI data traffic for one iSCSI link is carried on one TCP connection. Consequently, the aggregated bandwidth will be one Gbps for that iSCSI link.

Figure 24-31 provides a sample Ethernet PortChannel-based high availability iSCSI configuration.

Figure 24-31 Ethernet PortChannel-Based iSCSI High Availability

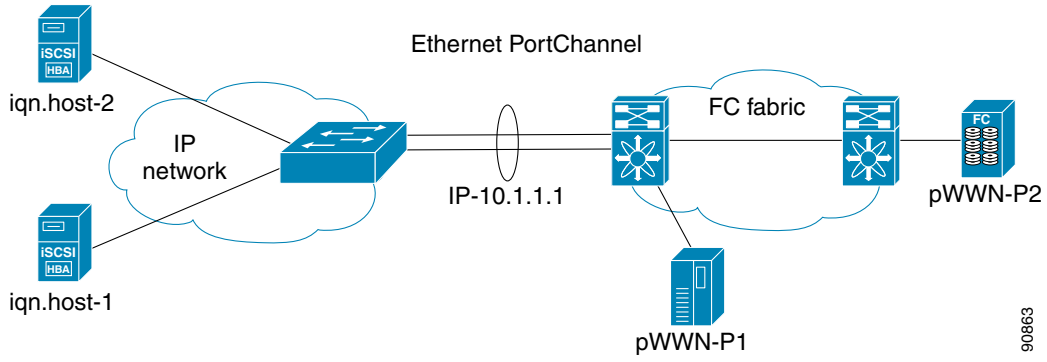

In Figure 24-31, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. The iSCSI session from the iSCSI host to the virtual iSCSI target (on the IPS port) uses one of the two physical interfaces (because an iSCSI session uses one TCP connection). When the Gigabit Ethernet 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

This section provides guidelines on iSCSI authentication possibilities, setup requirements, and sample scenarios.

This section does not specify the steps to enter or exit EXEC mode, configuration mode, or any submode. Be sure to verify the prompt before issuing any command.

Configuring Storage Name Services

Effective Release 1.3(1), the Internet Storage Name Service (iSNS) client feature is available in all switches in the Cisco MDS 9000 Family with IPS modules installed.

iSNS services allow your existing TCP/IP networks to function more effectively as storage area networks by automating the discover and management of iSCSI devices. To facilitate these functions, the iSNS client functionality registers iSCSI portals and all targets accessible through a particular interface with an external iSNS server.

Creating iSNS Profiles and Tagging Profiles

The iSNS client functionality on each interface (Gigabit Ethernet interface or subinterface or PortChannel) registers information with its configured iSNS server using an iSNS profile. This process is referred to as tagging an iSNS profile to an interface. Each iSNS profile keeps information about an iSNS server IP address. One profile can be tagged to one or more interfaces.

Once a profile is tagged to an interface, the MDS switch opens a TCP connection to the iSNS server IP address (using a well-known iSNS port number 3205) in the profile and registers network entity and portal objects. It goes through the FC name server database and configuration to find storage nodes to register with the server.

Statically-mapped virtual target is registered if the associated Fibre channel pWWN is present in the FC name server database and no access control configuration prevents it. A dynamically-mapped target is registered if the dynamic target importing is enabled.

A storage node is deregistered from the iSNS server when it becomes unavailable either because of configuration change (such as access control change or dynamic import disabling) or when the Fibre Channel storage port goes off-line. It will be registered again when the node is online.

When the iSNS client is unable to register/deregister objects with the iSNS server (e.g. as in unable to make a tcp connection to the iSNS server) it retries every minute to re-registers all iSNS objects for the affected interface(s) with the iSNS server.

Untagging a profile causes the network entity and portal to deregister from that interface.

Creating an iSNS Profile

To create an iSNS profile, follow these steps:

Step 1 Choose **End Devices > iSCSI** from the Fabric Manager Physical Attributes pane, or choose **IP > iSCSI** from Device Manager. You see the iSCSI tables in the Fabric Manager Information pane, or the Device Manager iSCSI dialog box.

Step 2 Click the **iSNS Profiles** tab. You see the iSCSI authentication configuration table.

Step 3 Set the Name and IP address of the iSNS server.

Step 4 Click the Fabric Manager **Apply Changes** icon (or click **Create** in Device Manager) to save these changes, or click **Undo Changes** in Fabric Manager (or click **Cancel** in Device Manager) to discard changes.

Modifying an iSNS Profile

To modify (tag) the iSNS profile for an interface, untag the interface from currently tagged iSNS profile and then tag to a new iSNS profile.

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

Step 1 Choose **Interfaces > Ethernet or iSCSI**. You see the Interfaces dialog box.

Step 2 Click the **General** tab. You see the General interface configuration table.

Step 3 Choose the iSNS ProfileName.

Step 4 Click **Apply** to save these changes, or click **Cancel** to discard changes.

Default IP Storage Settings

Table 24-1 lists the default settings for Gigabit Ethernet parameters.

Table 24-1 Default Gigabit Ethernet Parameters

Parameters	Default
IP MTU frame size	1500 bytes for all Ethernet ports

Table 24-2 lists the default settings for FCIP parameters.

Table 24-2 Default FCIP Parameters

Parameters	Default
TCP default port for FCIP	3225
minimum-retransmit-time	200 milliseconds.
keepalive-timeout	60 seconds.
max-retransmissions	4 retransmissions.
PMTU discovery	Enabled.
pmtu-enable reset-timeout	3600 seconds.
SACK	Enabled.
max-bandwidth	1G.
min-available-bandwidth	15 Mbps.
round-trip-time	1 ms.
buffer size	0 KB.
Control TCP and data connection	No packets are transmitted.
TCP congestion window monitoring	Enabled
Burst size	10KB.
TCP connection mode	active mode is enabled.
special-frame	Disabled.
FCIP timestamp	Disabled.
acceptable-diff range to accept packets	+ or - 1000 milliseconds.
B port keepalive responses	Disabled

Table 24-3 lists the default settings for iSCSI parameters.

Table 24-3 Default iSCSI Parameters

Parameters	Default
Number of TCP connections	One per iSCSI session.
Fibre Channel targets to iSCSI	Not imported.

Table 24-3 Default iSCSI Parameters (continued)

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 allowed to access a virtual target (unless explicitly configured).
iSCSI login authentication	CHAP or none authentication mechanism.
Ethernet PortChannel IP address usage	Source and destination IP addresses.

Using the IP Filter Wizard

Use the IP Filter Wizard to manage IP filters.

Step 1 From the Fabric Manager, click the **IP Filter** icon on the Fabric Manager toolbar.

Figure 24-32 IP Filter Wizard

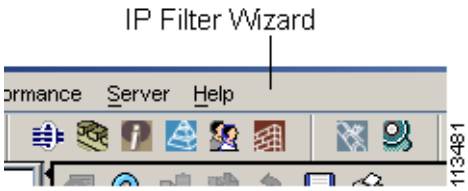

The IP Filter Wizard is displayed.

Step 2 Follow the prompts in the wizard to manage IP filters.

Creating IP Profiles

To create an IP profile, perform the following steps.

Step 1 From the Fabric Manager, choose **Security > IP Filter** from the menu tree. You see IP filter information on the Fabric Manager Information pane.

Step 2 Click the **Profiles** tab to see a list of profiles.

Step 3 Click the **Create Row** icon. You see the Create Profile dialog box.

Step 4 Choose the switches you want to include in the profile by checking the check box next to the switch's address.

Step 5 Enter a profile name in the Name field.

Step 6 Click **Create** to create the profile, or click **Close** to close the Create Profile dialog box without creating a profile.

You see the newly created profile in the list of profiles.

Step 7 To create additional profiles, repeat Steps 4 through 6. Otherwise, click **Close** to close the Create Profile dialog box.

Adding IP Filters to Profiles

To add an IP filter to a profile, perform the following steps.

Step 1 Choose **Security > IP Filter** from the Fabric Manager menu tree. You see IP Filter information on the Information pane.

Step 2 Click the **Profiles** tab to see a list of switches and associated profiles.

Step 3 Click on the IP address of the switch to which you want to add a filter. The Rules button becomes available.

Step 4 Click **Rules**. You see the IP Filter Edit dialog box.

Step 5 Click **Create Row**. You see the Create IP Filter dialog box.

Step 6 Complete the fields in the Create IP Filter dialog box.

Step 7 Click **Create** to create the filter, or click **Close** to close the Create IP Filter dialog box without creating a filter.

You see the newly created filter in the list of filters.

Step 8 Repeat Steps 6 and 7 to create additional filters, or click **Close** to close the Create IP Filter dialog box.

Step 9 Click **Apply Changes** to add the newly created filters to the profile.

Associating IP Profiles to Interfaces

To associate the profile to an interface, perform the following steps.

Step 1 Choose **Security > IP Filter** from the Fabric Manager menu tree. You see the IP Filter information in the Information pane.

Step 2 Click the **Interfaces** tab to see a list of interfaces and associated profiles.

Step 3 Click the **Create Row** icon. You see the Create Interface dialog box.

Step 4 Select the switches you want to include in the profile, by checking the check boxes next to the switch's address.

Step 5 Enter an interface name in the Name field.

Step 6 Choose the profile direction (either inbound or outbound).

Step 7 Enter the profile name in the Profile Name field. (This profile name must already have been created using the Create Profiles dialog. If not, no filters will be enabled until you to go to the Create Profiles dialog and create the profile.)

Step 8 Click **Create** to associate the profile, or click **Close** to close the Create Interfaces dialog box without associating a profile.

You see the newly associated profile in the list of profiles.

Step 9 Repeat Steps 4 through 8 to create additional associations, or click **Close** to close the Create Interfaces dialog box.

Deleting IP Profiles

To delete an IP profile, perform the following steps.

Step 1 Choose **Security > IP Filter** from the Fabric Manager menu tree. You see IP Filter information in the Information pane.

Step 2 Click the **Profiles** tab to see a list of switches, profile names and profile types.

Step 3 Click the row you want to delete. If you want to delete multiple rows, hold down the **Shift** key while clicking rows.

Step 4 Click the **Delete Row** icon to delete the profile.

Deleting IP Filters

To delete an IP filter, perform the following steps.

Step 1 Choose **Security > IP Filter** from the Fabric Manager menu tree. You see IP Filter information in the Information pane.

Step 2 Click the **Interfaces** tab to see a list of switches, filters and profile names.

Step 3 Click the row you want to delete. If you want to delete multiple rows, hold down the **Shift** key while clicking rows.

Step 4 Click the **Delete Row** icon to delete the filter from the profile

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.

This chapter provides configuration and messaging details on the Call Home feature.

This chapter contains the following topics:

- Call Home Features, page 25-2
- Call Home Configuration Process, page 25-2
- Cisco AutoNotify, page 25-2
- Assigning Contact Information, page 25-3
- Configuring Destination Profiles, page 25-3
- Configuring Alert Groups, page 25-3
- Configuring Message Levels, page 25-4
- Configuring E-Mail Options, page 25-4
- Enabling or Disabling Call Home, page 25-4
- Default Settings, page 25-4
- Event Triggers, page 25-4
- Call Home Message Severity Levels, page 25-7
- Message Contents, page 25-8
- Call Home Configuration Overview, page 25-11
- Configuring Call Home Attributes, page 25-12
- Configuring Call Home Destination Attributes, page 25-12
- Configuring Call Home E-Mail Addresses, page 25-12
- Configuring Call Home Alerts, page 25-13
- Configuring Call Home Profiles, page 25-13

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. Each profile may be predefined or user-defined.

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.
  - http://www.cisco.com/—The XML format enables communication with the Cisco Systems TAC group.
- Multiple concurrent message destinations. Up to 50 E-mail destination addresses are allowed for each format type.
- Message categories include system, environment, switching module hardware, supervisor module, hardware, inventory, and test.

Call Home Configuration Process

The actual configuration of Call Home depends on how you intend to use the feature. Some points to consider include:

To configure Call Home, follow these steps:

Step 1 Configure the Call Home function (see the "Call Home Configuration Overview" section on page 25-11).

Step 2 Assign contact information (see the "Assigning Contact Information" section on page 25-3).

Step 3 Configure destination profiles (see the "Configuring Destination Profiles" section on page 25-3).

Step 4 Enable or disable Call Home (see the "Enabling or Disabling Call Home" section on page 25-4).

Step 5 Test Call Home messages.

Cisco AutoNotify

For those who have service contracts directly with Cisco Systems, automatic case generation with the Technical Assistance Center is possible through registration with the AutoNotify service. AutoNotify provides fast time to resolution of system problems by providing a direct notification path to Cisco customer support.

To register, the following items are required:

- The SMARTnet contract number covering your MDS 9000 family switch.
- Your name, company address, your email address, and your CCO ID.
- 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 exact product number of your Cisco MDS 9000 Family switch. For example, some valid product numbers include DS-C6509 and DS-C9216-K9

To configure a Cisco MDS 9000 Family switch to use AutoNotify service, an XML destination profile must be configured to send messages to Cisco. Specific setup, activation, and email address information is found on the Cisco.com web site at:

The AutoNotify feature requires several Call Home parameters to be configured, including certain contact information, email server, and an XML destination profile as specified in the Service Activation document. The **contract-id**, **customer-id**, **site-id**, and **switch-priority** parameters are not required by the AutoNotify feature. They are only intended to be used as additional information by Cisco customers and service partners.

Assigning 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.

Configuring 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.

If you use the Cisco AutoNotify service, the XML destination profile is required. (See the "Configuring Destination Profiles" section on page 25-3)

- 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).

Configuring Alert Groups

The **alert-group** option allows you to select predefined types of Call Home alert notifications for destination profiles (predefined and user-defined). Destination profiles can be associated with multiple alert groups.

Configuring Message Levels

The **message-level** option 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 will not be sent. The urgency level ranges from 0 (lowest level of urgency) to 9 (highest level of urgency), and the default is 0 (all messages will be sent).

Configuring 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 ensure to configure the SMTP server address and port number for the Call Home functionality to work.

Enabling or Disabling Call Home

Once you have configured the contact information, you must enable the Call Home function.

Default Settings

Table 25-1 lists the default Call Home default settings.

Table 25-1 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

Event Triggers

This section discusses Call Home trigger events. Trigger events are divided into categories, with each category assigned commands to execute when the event occurs. The command output is included in the transmitted message. Table 25-2 lists the trigger events. lists event categories and command outputs.

Table 25-2 Event Triggers

Event	Alert Group	Event Name	Description	Severity 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		
Call Home	Supervisor Hardware and CISCO_TAC	INBAND_FAILURE	Failure of inband communications path	7

Table 25-2 Event Triggers (continued)

Event	Alert Group	Event Name	Description	Severity Level
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 that is not licensed (Release 1.3.x), and will be turned off after grace period expiration	6	
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

Table 25-3 Event Categories and Command Outputs

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**

Table 25-3 Event Categories and Command Outputs (continued)

Event Category	Description	Executed Commands
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 Severity Levels

This section discusses the severity levels for a Call Home message when using one or more switches in the Cisco MDS 9000 Family. Severity levels are preassigned per event type.

Call Home severity levels are not the same as system message logging severity levels. Severity levels range from 0 to 9, with 9 having the highest urgency. Each severity level has keywords as listed in Table 25-4.

Table 25-4 Severity Levels

Severity Level	Keyword	Description
9	Catastrophic	Network wide catastrophic failure.
8	Disaster	Significant network impact.
7	Fatal	System is unusable.
6	Critical	Critical conditions, immediate attention needed.
5	Major	Major conditions.
4	Minor	Minor conditions.
3	Warning	Warning conditions.
2	Notification	Basic notification and informational messages. Possibly independently insignificant.
1	Normal	Normal event signifying return to normal state.
0	Debugging	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 25-5 describes the short text formatting option for all message types.

Table 25-5 Short Text Formatting Option

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 syslog message

Table 25-6 displays the information contained in plain text and XML messages.

Table 25-6 Plain Text and XML Messages

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-DD T HH: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.	/mml/header/name
Message type	Specifically “Call Home”.	/mml/header/type
Message group	Specifically “reactive”.	/mml/header/group
Severity level	Severity level of message.	/mml/header/level
Source ID	Product type for routing.	/mml/header/source

Table 25-6 Plain Text and XML Messages (continued)

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
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 serial ID as a chassis serial number· • Serial number as identified by the Sid field. Example: 'DS-C9000@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

Table 25-6 Plain Text and XML Messages (continued)

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
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 serial ID as a chassis serial number· • Serial number as identified by the Sid field. Example: 'DS-C9000@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
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhoneNumber
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

Table 25-6 Plain Text and XML Messages (continued)

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
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	Exact command that was run. For example, **show running-config** 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.	/mml/attachments/attachment/at data

Call Home Configuration Overview

When configuring Call Home, keep the following points in mind:

- You must configure at least one E-mail server and at least one destination profile. The destination profile(s) used depends on whether the notification is sent to a pager, e-mail, or automated service such as Cisco AutoNotify.
- You must configure the contact name (SNMP server contact), phone, and street address information before enabling Call Home.
- The Cisco MDS 9000 switch must have IP connectivity to an E-mail server.
- To use Cisco AutoNotify you must obtain an active service contract for the device.

To configure Call Home, use the different tabs on the Call Home dialog box, as summarized below:

Step 1 Assign contact information and enable the Call Home feature using the General tab (see the "Configuring Call Home Attributes" section on page 25-12). The Call Home feature is not enabled by default, and you must enter an e-mail address that identifies the source of Call Home notifications.

Step 2 Configure the destination e-mail addresses for Call Home notifications using the Destinations tab (see the "Configuring Call Home Destination Attributes" section on page 25-12). You can identify one more more e-mail addresses that will receive Call Home notifications.

Step 3 Identify your SMTP server using the E-mail Setup tab ("Configuring Call Home E-Mail Addresses" section on page 25-12). You need to identify a message server to which your switch has access. This message server will forward the Call Home notifications to the destinations.

Step 4 Test Call Home by sending a test message using the Alerts tab (see the "Configuring Call Home Alerts" section on page 25-13). You should test the Call Home feature to make sure it works.

Configuring Call Home Attributes

To assign contact information and enable the Call Home feature from the Fabric Manager, choose Events > Call Home on the menu tree and click the General tab. The Information pane from the Fabric Manager displays Call Home information for multiple switches.

To assign contact information and enable the Call Home feature from the Device Manager, choose Call Home from the Events menu and click the General tab. The Call Home Events dialog box with the General tab selected from the Device Manager displays Call Home attributes for a single switch.

Configuring Call Home Destination Attributes

To configure the destination e-mail addresses for Call Home notifications from the Fabric Manager, choose Events > Call Home on the menu tree and click the Destination tab. The Information pane from the Fabric Manager displays Call Home information for multiple switches.

To configure the destination e-mail addresses from the Device Manager, choose Call Home from the Events menu and click the Destination tab. The dialog box from the Device Manager displays Call Home attributes for a single switch.

To create a new Call Home destination, follow these steps:

Step 1 Click **Create** in the Device Manager dialog box, or click **Create Row** on the Fabric Manager toolbar.

You see the Device Manager Create Call Home Destination dialog box.

From the Fabric Manager, you can choose one or more switches to which the configuration applies.

Step 2 Choose the profile name from the pull-down list.

Step 3 Enter a number identifier for the destination.

Step 4 Enter the e-mail address for the destination.

Step 5 Click **Create**.

Configuring Call Home E-Mail Addresses

To identify your SMTP server from the Fabric Manager, choose Events > Call Home on the menu tree and click the Email Setup tab. The Information pane from the Fabric Manager displays Call Home information for multiple switches.

To identify your SMTP server from the Device Manager, choose Call Home from the Events menu and click the Email Setup tab. The Call Home dialog box from the Device Manager displays Call Home attributes for a single switch.

Configure the e-mail setup attributes for the Call Home features.

Configuring Call Home Alerts

To test Call Home from the Fabric Manager, choose Events > Call Home the menu tree and click the Alerts tab. The Information pane from the Fabric Manager displays Call Home information for multiple switches.

To test Call Home from the Device Manager, choose Call Home from the Events menu and click the Alerts tab. The dialog box with the Alerts tab selected from the Device Manager displays Call Home attributes for a single switch.

Configure the alert attributes for the Call Home feature.

Configuring Call Home Profiles

To configure Call Home attributes from the Fabric Manager, choose Events > Call Home on the menu tree and click the Profiles tab. The Information pane from the Fabric Manager displays Call Home information for multiple switches.

To configure Call Home attributes from the Device Manager, choose Call Home from the Events menu and click the Profiles tab. The dialog box with the Alerts tab selected from the Device Manager displays Call Home attributes for a single switch.

Configure the profile attributes for the Call Home feature.

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.

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 contains the following topics:

- About fcdomain Phases, page 26-2
- Restarting the Domain, page 26-3
- Performing a Domain Restart, page 26-3
- Configuring the Domain, page 26-3
- Specifying a Preferred or Static Domain ID, page 26-6
- Setting Switch Priority, page 26-6
- Configuring Allowed Domain ID Lists, page 26-6
- Merging Stable Fabrics, page 26-7
- Assigning Contiguous Domains, page 26-7
- Disabling the fcdomain Feature, page 26-7
- Setting the Fabric Name, page 26-8
- Stopping Incoming RCFs, page 26-8
- Configuring Persistent FC IDs, page 26-8
- Enabling Persistent FC IDs, page 26-10
- Purging Persistent FC IDs, page 26-11
- Default Settings, page 26-12

About fcdomain Phases

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.

Figure 26-1 Sample fcdomain Configuration

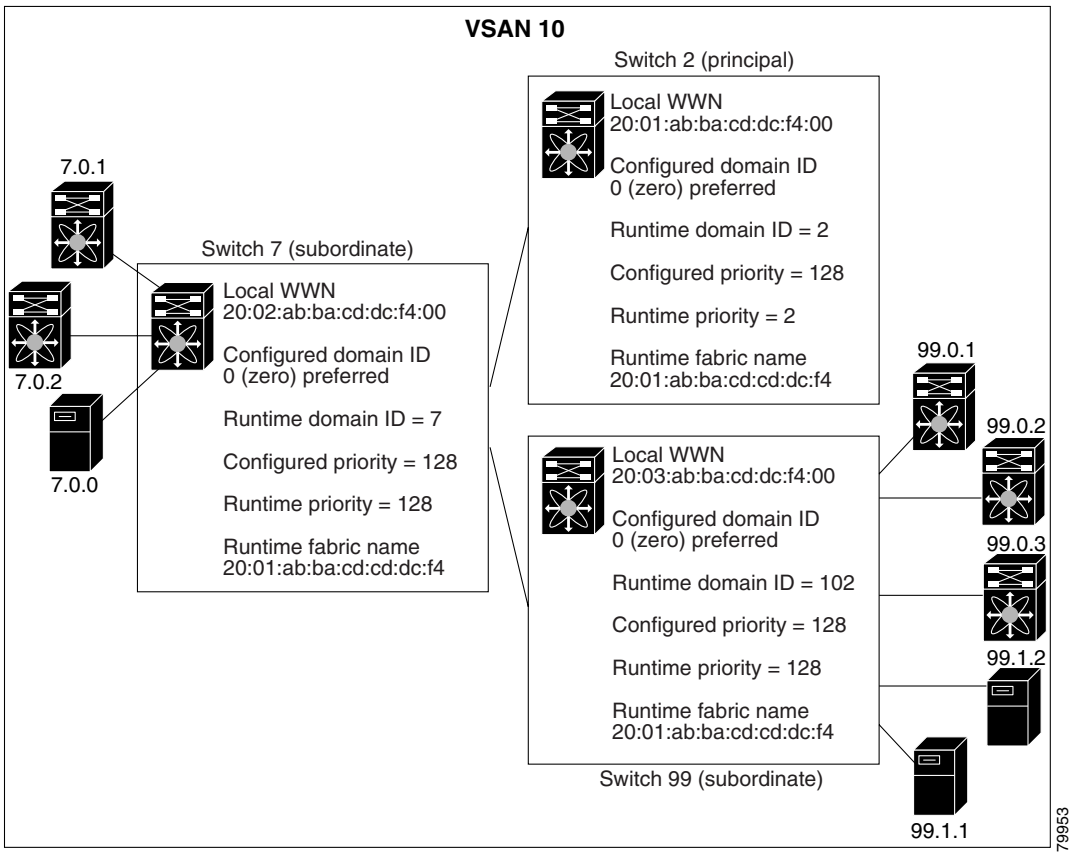

Note Domain IDs and VSAN values used in all procedures are only provided as examples. Be sure to use IDs and values that apply to your configuration.

Restarting the Domain

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. If you perform a nondisruptive restart, build fabric (BF) frames are sent to other switches in the fabric.

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 will change to take on the static domain ID after the next restart.

Tip If a VSAN is in **interop** mode, you cannot restart the fcdomain for that VSAN disruptively.

Performing a Domain Restart

To restart the fabric disruptively or nondisruptively, follow these steps:

Step 1 Open Device Manager on a switch in the domain you want to restart.

Step 2 Choose **FC > Domain Manager**. You see the Domain Manager dialog box.

Step 3 Click in the Restart column for the VSAN ID for which you want to restart the domain. You see a drop-down list of options.

Step 4 Choose **disruptive** for a disruptive restart, or **Non-disruptive** for a non-disruptive restart.

Step 5 Click **Apply** to restart the domain.

Configuring the Domain

The configured domain ID can be *preferred* or *static*. By default, the configured domain is 0 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.

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 26-2 Configuration Process Using the "preferred" Option

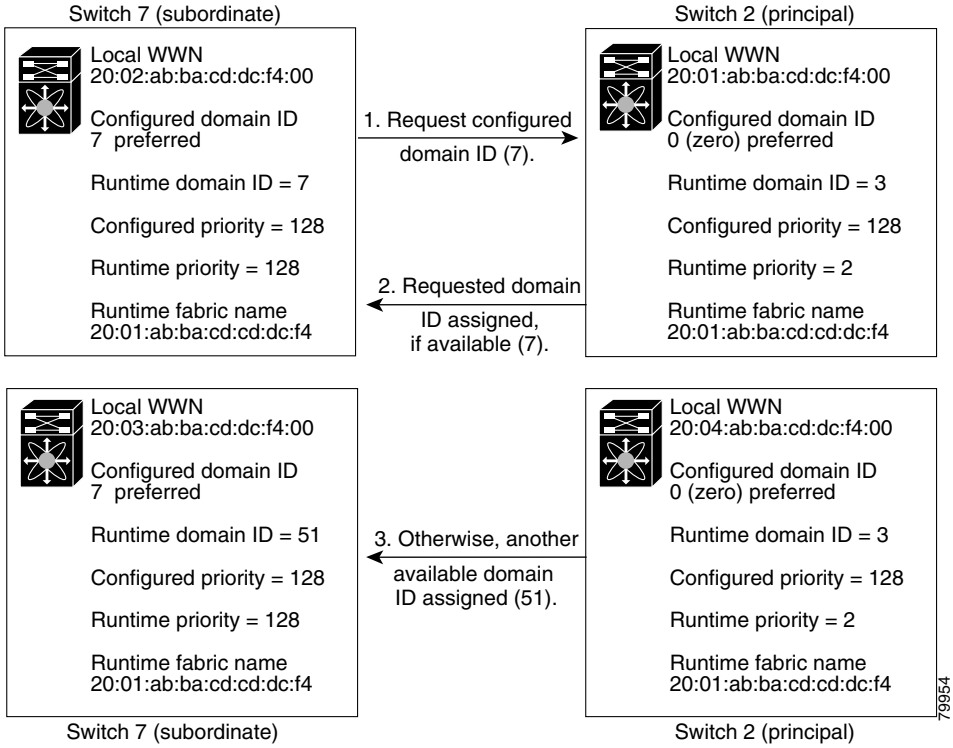

The behavior for a subordinate switch changes based on the allowed domain ID lists, on the configured domain ID, and on the domain ID that the principal switch has assigned to the requesting switch:

- 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 option 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 option is **preferred**, the local switch accepts the domain ID assigned by the principal switch and the assigned domain ID becomes the runtime domain ID.

Caution You must restart the domain if you want to apply the configured domain changes to the runtime domain.

Configuring Domain Attributes

From this dialog box you can specify a fabric name for fabric logins on the VSAN and set the priority for the switch used in the principal switch selection process.

Configure the principal attributes for the domain.

To manage domain attributes from the Fabric Manager, choose **FC > Domain Manager** on the menu tree and click the **Configuration** tab. The Information pane from the Fabric Manager lets you manage domain attributes for multiple switches.

To manage domain attributes from the Device Manager, choose **Domain Manager** from the FC menu and click the **Configuration** tab. The Device Manager dialog box displays domain attributes for a single switch.

Configure the attributes for the domain.

Managing Running Attributes for Domains

To view running domain attributes from the Fabric Manager, choose **FC > Domain Manager** on the menu tree and click the **Running** tab. The Information pane from the Fabric Manager displays domain attributes for multiple switches.

To view running domain attributes from the Device Manager, choose **Domain Manager** from the FC menu and click the **Running** tab. The Domain Manager dialog box, with the Running tab selected, displays domain attributes for a single switch.

Viewing Domain Information

To view domain information from the Device Manager, choose **Domain Manager** from the FC menu and click the **Domains** tab. The dialog box displays domain information for a single switch

Viewing Domain Manager Statistics

To monitor domain manager statistics from the Fabric Manager, choose **FC > Domain Manager** on the menu tree and click the **Statistics** tab. The Information pane from the Fabric Manager displays domain statistics for multiple switches.

To monitor domain manager statistics from the Device Manager, choose **Domain Manager** from the FC menu and click the **Statistics** tab. The Domain Manager dialog box, with the Statistics tab selected, displays domain statistics for a single switch.

Configuring Domain Interfaces

To configure domain interfaces from the Fabric Manager, choose **FC > Domain Manager** on the menu tree and click the **Interfaces** tab. The Information pane from the Fabric Manager displays domain interfaces for multiple switches.

To configure domain interfaces from the Device Manager, choose **Domain Manager** from the FC menu and click the **Interfaces** tab. The Domain Manager dialog box, with the Interfaces tab selected, displays domain interfaces for a single switch.

Configure the attributes for domain interfaces.

Viewing Domain Areas

To monitor domain areas from the Device Manager, choose **Domain Manager** from the FC menu and click the **Areas** tab. The Domain Manager dialog box, with the Areas tab selected, displays domain areas for a single switch.

Viewing Domain Area Ports

To monitor area ports for domains from the Device Manager, choose **Domain Manager** from the FC menu and click the **Area Ports** tab. The Domain Manager dialog box, with the Area Ports tab, displays area ports for domains for a single switch.

Specifying a Preferred or Static 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 to runtime after a disruptive or nondisruptive restart, the **preferred** option is applied to runtime only after a disruptive restart.

Tip When the FICON feature is enabled in a given VSAN, the domain ID for that VSAN will remain in the static state. You can change the static ID value but you cannot change it to the preferred option.

Setting 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.

Configuring 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 ensures that the domain requested by any switch in the fabric is specified in the allowed list.

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.

Merging Stable Fabrics

By default, the **auto-reconfigure** option is disabled. When you join two switches belonging to two different stable fabrics that have overlapping domains, the following cases apply:

- If the **auto-reconfigure** option is enabled on both switches, a disruptive reconfiguration phase is started.
- If the **auto-reconfigure** option is disabled on either or both switches, the links between the two switches become isolated.

The **auto-reconfigure** 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. However, if you enable the **auto-reconfigure** option on both switches before connecting the fabric, a disruptive reconfiguration (RCF) occurs. A disruptive reconfiguration may affect data traffic. You can nondiruptively reconfigure the fcdomain by changing the configured domains on the overlapping links and getting rid of the domain overlap.

Assigning Contiguous Domains

By default, the **contiguous-allocation** option is disabled. When the subordinate switches request the principal switch for two or more domains and the domains are not contiguous, the following cases apply:

- If the **contiguous-allocation** option is enabled in the principal switch, the principal switch locates contiguous domains and assigns them to the subordinate switches.
- If the **contiguous-allocation** option is disabled in the principal switch, the principal switch assigns the available domains to the subordinate switches.

The **contiguous-allocation** option takes immediate effect at runtime; you do not need to restart the fcdomain.

Disabling the fcdomain Feature

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.

Setting the Fabric Name

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 WWN.

The fabric name is applied to runtime through a disruptive restart when the fcdomain is configured as disabled.

Stopping Incoming RCFs

The **rcf-reject** option is configured 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.

Configuring Persistent FC IDs

By default, the persistent FC_ID feature is disabled. When a N/NL-port logs into a switch, and gets assigned an FC ID, the WWN of the requesting N/NL-port and the assigned FC ID is retained and stored in a volatile cache (the content is lost after a reboot).

If the persistent FC ID feature is disabled, binding of the FC ID to the WWN is preserved on a best effort basis.

For example, after the disconnection of one N-Port from the switch, if its FC ID is requested by another device, the request is granted and the initial association WWN FC ID is released. Also, if the 4K entries of the volatile cache used to store the WWN-to- FC ID binding get completely filled up, a new (more recent) entry will overwrite the oldest one, losing the corresponding binding WWN to FC ID.

The behavior is different for an N-Port than for an NL-Port:

- N-ports should receive the same FC IDs if unplugged and plugged back in any port of the same switch (as long as it belongs to the same VSAN)
- NL-port should receive the same FC IDs only if connected back to the same interface on the switch it was connected originally.

The assigned FC IDs in a fcdomain can be activated to remain persistent, even after a reboot. This ensures that an attached N Port receives the same FC ID after a switch reboot. If you enable this feature, the following apply:

- The currently “in-use” FC IDs in the fcdomain will be 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.

To configure persistent FC IDs from the Fabric Manager, choose **FC > Domain Manager** on the menu tree and click the **Persistent FCIDs** tab. The Information pane from the Fabric Manager displays persistent FC IDs for multiple switches.

To configure persistent FC IDs from the Device Manager, choose **Domain Manager** from the FC menu and click the **Persistent FCIDs** tab. The Domain Manager dialog box, with the Persistent FCIDs tab selected, displays persistent FC IDs for a single switch.

Configure the attributes for persistent FC IDs.

Before you can create persistent FC IDs, you must:

- Configure a static domain ID in that VSAN
- Ascertain that the static configured domain and the runtime domain are the same. You can verify this using the show fcdomain command. For information about using the command line interface (CLI), refer to the Cisco 9000 Family Configuration Guide.

Note If you connect to the switch from an AIX or HP-UX host, be sure to create the persistent FC ID in the VSAN that connects these hosts.

Note Persistent FC IDs with loop-attache devices (FL ports) need to remain connected to the same port in which they were configured.

Creating a Persistent FC ID

To create a new persistent FC ID, follow these steps:

Step 1 Click **Create**. You see the Create Domain Manager Persistent FCIDs dialog box.

Step 2 Enter the VSAN ID.

Step 3 Enter the WWN.

Step 4 Enter the FC ID.

Step 5 Choose the mask. The mask is the number of FC IDs which are assigned either statically or dynamically for this WWN on this VSAN. Possible values are *Single*, meaning just one FCI ID is assigned, or *Area*, meaning all of the FC IDs in the area that is specified are assigned.

Step 6 Choose the assignment. The assignment is the type of persistency of this FC ID (static or dynamic).

Step 7 Click **Create** to create the persistent FC ID, or click **Close** to return to the Domain Manager without creating the FC ID.

Deleting a Persistent FC ID

To delete a persistent FC ID, follow these steps:

Step 1 Choose the persistent FC ID you want to delete. This enables the Delete button.

Step 2 Click **Delete** to delete the FC ID.

Configuring Persistent FC IDs Manually

Once 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.
- Do not replace an FC ID that is already configured in another WWN. If you want to use a previously-configured WWN, first delete the configured WWN before proceeding with this procedure.

Note You cannot configure persistent FC IDs in FICON-enabled VSANs.

Configuring Unique Area FC IDs for Some HBAs

Note This section does not apply if either the HBA port or the storage port is connected to different switches.

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 area can be anything other than 77. The HBA port FC ID must be manually configured to be different from the storage port 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 FCC 1/10 in the same switch.

Enabling Persistent FC IDs

Persistent FC IDs are disabled by default. You can enable this option globally or for each VSAN. If you choose to enable it globally, you can do so at any time using the initial setup routine or the setup command. When you enable this option globally, the switch remains in this state until you change the global configuration.

Note If you enable this option during the initial switch setup, this option will be automatically enabled in all configured VSANs. If you enable this option at a later stage, this option will be automatically enabled in all VSANs configured after that stage. VSANs configured before that stage will remain unchanged.

When a N or NL port logs into a Cisco MDS 9000 Family switch, it is assigned a FC ID. By default, the persistent FC ID feature is disabled. If this feature is disabled, the following consequences apply:

- A 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 it was originally connected.

The assigned FC IDs in a fcdomain can be enabled to remain persistent even after a reboot. This ensures that an attached N port receives the same FC IDs after a reboot. If you enable this feature, 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 Persistent FC IDs with loop-attached devices (FL ports) need to remain connected to the same port in which they were configured.

A persistent FC ID assigned to an F port can be moved across interfaces and can continue to maintain the same persistent FC ID.

Purging Persistent FC IDs

Persistent FC IDs can be purged selectively. Static entries and FC IDs currently in use cannot be deleted. Table 26-1 identifies the FC ID entries that can be deleted.

Table 26-1 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

Default Settings

Table 26-2 lists the default settings for all fcdomain parameters.

Table 26-2 Default fcdomain Parameters

Parameters	Default
fcdomain feature	Enabled.
Configured domain ID	0 (zero).
Configured domain option	Preferred.
auto-reconfigure option	Disabled.
contiguous-allocation option	Disabled.
Priority	128.
Allowed list	1 to 239.
Fabric-name	20:01:00:05:30:00:28:df.
rcf-reject	Disabled.
Persistent FC ID	Disabled (globally configurable).

Configuring Traffic Management

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.

This chapter contains the following topics:

- FCC, page 27-1
- QoS, page 27-2
- Control Traffic, page 27-3
- Data Traffic, page 27-3
- Ingress Port Rate Limiting, page 27-6
- Default Settings, page 27-6

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.

Figure 27-1 FCC Mechanisms

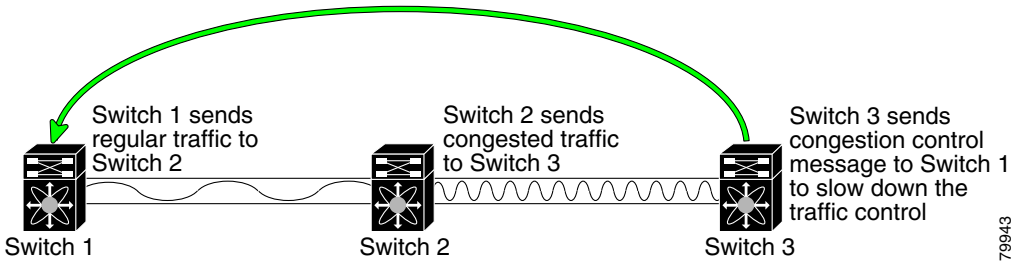

Edge quench congestion control provides feedback to the source about the rate at which frames should be injected into the network (frame intervals).

FCC Process

When a node in the network detects a 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 quest 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 quest frames. However, only the edge switch processes edge quest frames.

Enabling FCC

By default, the FCC protocol is disabled. FCC can only be enabled for the entire switch.

If you enable FCC, be sure to enable it in all switches in the fabric.

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:

- Control Traffic
- DataTraffic

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 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.

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.

We do not recommended disabling this feature as all critical control traffic will automatically be assigned the lowest priority once you issue this command. You can view the current state of the QoS configuration for critical control traffic using the **show qos statistics** command.

Data Traffic

Transaction processing, a low volume, latency sensitive application, requires quick access to requested information. Backup processing requires high bandwidth but is not sensitive to latency. In a network that does not support service differentiation, all traffic is treated identically; they experience similar latency and get similar bandwidths. The QoS feature in all switches in the Cisco MDS 9000 Family provides these guarantees from SAN-OS Release 1.3(x).

Prior versions of the SAN-OS software only differentiated traffic priority based on control traffic. SAN-OS Release 1.3(x) enables you to take full advantage of the QoS capabilities. Data traffic can now be prioritized in four distinct levels of service differentiation: low, medium, high, or absolute priority. You can apply QoS to ensure that Fibre Channel data traffic for your latency-sensitive applications receive higher priority over throughput-intensive applications like data warehousing.

Figure 27-2 Prioritizing Data Traffic

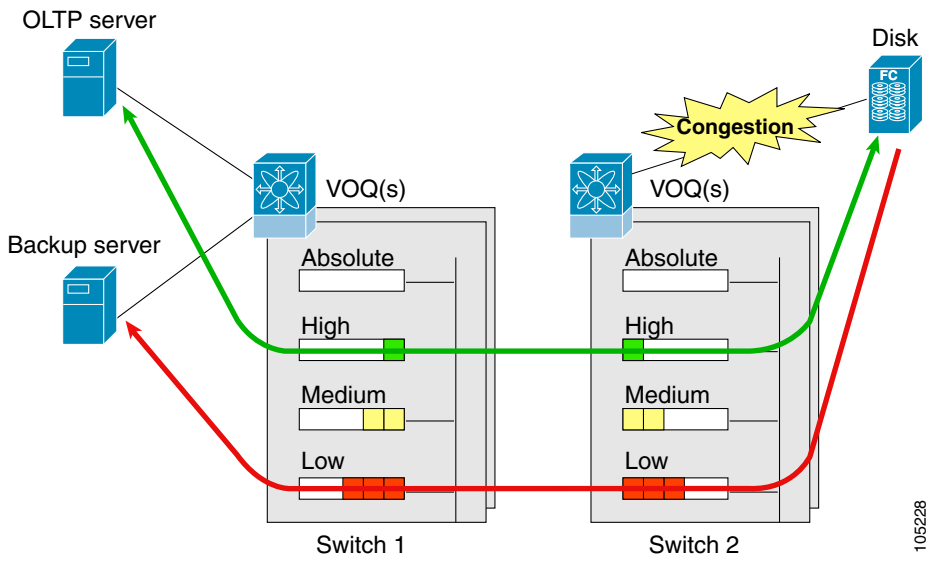

In Figure 27-2, the OLTP traffic arriving at Switch 1 is marked with a **High** priority level of through classification (class map) and marking (policy map). Similarly, the backup traffic is marked with a **Low** priority level. The traffic is sent to the corresponding priority queue within a Virtual 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 the high priority queue is not congested. The scheduler assigns it priority over the backup traffic in the low priority queue.

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.

To achieve this traffic differentiation, be sure to enable FCC.

Configuring Data Traffic

To configure QoS, 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.

Enabling QoS for Data Traffic

By default, the QoS data traffic feature is disabled for data traffic. To configure QoS for data traffic, you must first enable the data traffic feature in the switch.

QoS is supported in interoperability mode. The effectiveness of the setting depends on the location of MDS switches in the fabric relative to the location of the source or destination of the prioritized devices.

Creating Class Maps

Use the **class-map** option to create and define a traffic class with match criteria to identify traffic belonging to that class. Define each match criterion with one match statement from the class map configuration (switch(config-cmap)) mode. 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—Use the **source-wwn** option to specify the source WWN or the **destination-wwn** option to specify the destination WWN.
- Fibre Channel ID (FC ID)—Use the **source-address** option to specify the source ID (SID) or the **destination-address** option to specify the destination ID (DID). The possible values for mask are:
  - FFFFFF—The entire FCID is used. This is the default.
  - FFFF00—Only domain and area FCID is used.
  - FF0000—Only domain FCID is used.

  A source-address or destination-address of 0x000000 is not allowed.
- Source interface—Use the **input-interface** option to specify the ingress interface.

The order of entries to be matched within a class map is not significant.

Defining Service Policies

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.

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. See http://www.cisco.com/warp/public/105/dscpvalues.html#dscpandassuredforwardingclasses for further information on implementing QoS DSCP values.

Use the **policy-map** option to specify the class of service. The policy map name is restricted to 63 alphanumeric characters.

Class-maps are processed in the order in which they are configured in each policy-map.

Applying a Service Policy

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 will not be enforced. You can only apply one policy map to a VSAN.

You can apply the same policy to a range of VSANs.

Scheduling Traffic

The SAN-OS software supports four scheduling queues:

- Strict-priority queues are queues that are serviced in preference to other queues. A strict-priority queue 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 queues:
  - Use **dwrr-q high** option to schedule high priority traffic.
  - Use **dwrr-q medium** option to schedule medium priority traffic.
  - Use **dwrr-q low** option to schedule low priority traffic.

Use the **qos dwrr-q** command to associate a weight with a DWRR queue.

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 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).

Ingress Port Rate Limiting

A port rate limiting feature is available in SAN-OS 1.3(x). This feature helps control the bandwidth for individual FC ports. Port rate limiting is also referred to as ingress rate limiting because it controls ingress traffic into a FC 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.

Port rate limiting can only be configured in switches in the Cisco MDS 9100 Series.

This command can only be configured if the following conditions hold true:

- The QoS feature is enabled using the **qos enable** command.
- The command is issued in a Cisco MDS 9100 series switch.

The rate limit ranges form 1 to 100% and the default is 100%.

Default Settings

Table 27-1 lists the default settings for FCC, QoS, and rate limiting features:

Table 27-1 Default FCC, QoS, and Rate Limiting Settings

Parameters	Default
FCC protocol	Disabled.
QoS control traffic	Enabled.
QoS data traffic	Disabled.
Rate limit	100%

Configuring System Message Logging

This chapter describes how to configure system message logging on the Cisco MDS 9000 Family switches.

This chapter contains the following topics:

- About System Message Logging, page 28-1
- Configuring System Message Logging, page 28-3
- Default Settings, page 28-6
- About SNMP Events, page 28-7
- About RMON Facilities, page 28-8

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 of 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. Messages are time-stamped to enhance real-time debugging and management.

You can access logged system messages using the CLI or by saving them to a properly configured syslog server. The switch software saves syslog messages in a file that can be configured to save up to 4 MB. You can monitor system messages remotely by accessing the switch through Telnet, SSH, or the console port, or by viewing the logs on a syslog server.

When the switch first initializes, the network is not connected until initialization completes. Therefore, messages are not redirected to a syslog 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 28-1 describes the facilities supported by the system message logs.

Table 28-1 Facilities Supported by the System Message Logs

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 syslog messages	Standard
sysmgr	System manager	Cisco MDS 9000 Family specific
tlport	TL port	Cisco MDS 9000 Family specific

Table 28-1 Facilities Supported by the System Message Logs (continued)

Facility Keyword	Description	Standard or Cisco MDS Specific
user	User process	Standard
uucp	Unix-to-Unix copy system	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 syslog	Cisco MDS 9000 Family specific
vshd	vshd	Cisco MDS 9000 Family specific
wwn	WWN manager	Cisco MDS 9000 Family specific
xbar	Xbar syslog	Cisco MDS 9000 Family specific
zone	Zone server	Cisco MDS 9000 Family specific

Table 28-2 describes the severity levels supported by the system message logs.

Table 28-2 Error Message Severity Levels

Level Keyword	Level	Description	Syslog 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

Refer to the *Cisco MDS 9000 Family System Messages Guide* for details on the error log message format.

Configuring System Message Logging

System logging messages are sent to the console based on the default (or configured) logging facility and severity values.

Enabling Message Logging

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.

Configuring 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. T o increase the logging level (above critical), you must change the console baud speed to 38400 baud.

Configuring 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.

Configuring 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. You can rename this file. The file name can have up to 200 characters and the file size ranges from 4096 bytes to 4194304 bytes.

The configured log file is saved in the /var/log/external directory. The location of the log file cannot be changed.

You can display the log file and copy the logfile to a different location.

Configuring the Syslog Daemon

To send log messages to a UNIX syslog server, you must configure the syslog daemon on a UNIX server. Log in as root, and perform these steps:

Step 1 Add the following line to the file /etc/syslog.conf

```
local1.debug /var/log/ myfile.log
```

Be sure to add five tab characters between**local1.debug** and **/var/log/**myfile**.log**. Refer to 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 syslog daemon reads the new changes by entering this command:

```
$ kill -HUP ~cat /etc/syslog.pid~
```

Outgoing Syslog Server Logging Facilities

All syslog 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 syslog 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 28-3.

Table 28-3 Internal 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 syslog messages	Standard
user	User process	Standard
uucp	Unix-to-Unix copy system	Standard

Configuring Syslog Servers

To configure syslog servers, follow these steps:

Step 1 Choose **Events > Syslog** from the Fabric Manager menu tree, and then click the **Servers** tab. The Information pane displays syslog information for multiple switches.

From the Device Manager, choose **Events > Syslog**, and then click the **Servers** tab. The Syslog dialog box with the Servers tab selected displays syslog information for a single switch.

Step 2 Configure the server attributes for the syslog.

Step 3 To add a syslog server, click **Create**. You see the Create Syslog Server dialog box.

Step 4 Complete the fields on this dialog box and click **OK**.

Configuring Syslog Attributes

To configure syslog attributes, follow these steps:

Step 1 From the Fabric Manager, choose **Events > Syslog** on the menu tree and click the **General** tab. The Information pane displays syslog information for multiple switches.

From the Device Manager, choose **Events > Syslog** and click the **General** tab. The General tab of the Syslog dialog box displays syslog information for a single switch.

Step 2 Configure the general attributes for the syslog.

Configuring Syslog Priorities

To configure syslog priorities, follow these steps:

Step 1 From the Fabric Manager, choose **Events > Syslog** on the menu tree, and then click the **Priorities** tab. The Information pane displays syslog information for multiple switches.

From the Device Manager, choose **Events > Syslog**, and then click the **Priorities** tab. The Syslog dialog box with the Servers tab selected displays syslog information for a single switch.

Step 2 Configure the priorities for the syslog.

Default Settings

Table 28-4 lists the default settings for system message logging.

Table 28-4 Default System Message Log Setting

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 (can be changed to any name with up to 200 characters).
Logging server	Disabled.
Syslog server IP address	Non configured.
No. of servers	3 servers.
Server facility	Local 7.

About SNMP Events

SNMP is an application layer protocol that facilitates the exchange of management information between network devices. Cisco MDS 9000 Family switches, like other SNMP-enabled devices, send events (traps and informs) to configurable destinations, called trap receivers in SNMPv2.

Viewing the Events Log

To view the events log from the Device Manager, choose **SNMP Log** from the Events menu. The Events Log dialog box displays a log of events for a single switch

To manage the SNMP log, choose **SNMP Log** from the Events menu and click the **Controls** tab. The Controls tab provides summary statistics about the SNMP log and allows you to change the default settings for the log.

Caution Changing these values from different Fabric Manager workstations at the same time may cause unpredictable results.

Configuring Event Destinations

To configure event destinations, follow these steps:

Step 1 Choose **Events > Notifications/Traps** from the Fabric Manager the menu tree and click the **Destinations** tab, or choose **Events > Destinations** from the Device Manager.

The Fabric Manager Information pane from the Fabric Manager shows event destination information for multiple switches. The Device Manager dialog shows event destinations for a single switch.

Step 2 To create an event destination, click **Create** in the Device Manager dialog box or click the **Create Row** icon on the Fabric Manager toolbar. You see the Create Event Destinations dialog box. (The dialog box from the Fabric Manager allows you to choose a switch.)

Step 3 Complete the fields and click **Apply** to create the event destination, or click **OK** to create the destination and close the dialog box.

Configuring Event Security

Caution This is an advanced function that should only be used by administrators having experience with SNMPv3.

To configure event security from the Fabric Manager, choose **Events > Notifications/Traps** on the menu tree, and click the **Security** tab.

To configure event security from the Device Manager, choose **Destinations** from the Events menu and click the **Security** tab.

The Information pane from the Fabric Manager displays event security information for multiple switches. The dialog box from Device Manager displays event security for a single switch.

Configuring Event Filters

To configure event filters from the Fabric Manager, choose **Events > Filters** on the menu tree, and click the **FC or Other** tab.

To configure event filters from the Device Manager, choose **Filters** from the Events menu.

The Event Filters dialog box displays event filters for a single switch. The Information pane in Fabric Manager displays two different views, which list the same event filters for multiple switches, in different order.

To configure event filters, check the check box next to the appropriate filter name.

About RMON Facilities

Remote Monitoring (RMON) allows you to specify thresholds and monitor alarms on SNMP variables.

Enabling RMON Alarms by Port

To enable alarm notifications by port from the Device Manager, choose **Events > Threshold Manager** and click the **Ports** tab.

To configure an RMON alarm for one or more ports, follow these steps:

Step 1 Click the **Selected** radio button.

Step 2 Click the button to the right of the Selected field to display all ports.

Step 3 Choose the ports you want to monitor.

Step 4 Click **OK** to accept the selection.

Alternatively, click the appropriate radio button to select ports by type (All ports, xE ports, or Fx port).

Step 5 Check the check box for each variable that you want to monitor.

Step 6 Enter the threshold value in the Value column.

Step 7 Enter the sampling period in seconds.

Step 8 Choose one of the following severity levels to assign to the alarm:

- Fatal
- Warning
- Critical
- Error
- Information

Step 9 Click **Create**.

Step 10 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.

Enabling RMON Alarms for VSANs

To manage RMON alarm service attributes for selected VSANs from the Device Manager, choose **Events > Threshold Manager** and click the **Services** tab. You see the Threshold Manager dialog box with the Services tab selected.

To enable an RMON alarm for one or more VSANs, follow these steps:

Step 1 Enter one or more VSANs to monitor in the VSAN Id(s) field.

Step 2 Check the check box for each variable that you want to monitor.

Step 3 Enter the threshold value in the Value column.

Step 4 Enter the sampling period in seconds.

Step 5 Select a severity level to assign to the alarm:

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.

Enabling RMON Alarms for Physical Components

To configure RMON alarm physical attributes from the Device Manager, choose **Events > Threshold Manager** and click the **Physical** tab. You see the Create RMON Alarms dialog box with the Physical tab selected.

To configure an RMON alarm for a physical component, follow these steps:

Step 1 Check the check box for each variable that you want to monitor.

Step 2 Enter the threshold value in the Value column.

Step 3 Enter the sampling period in seconds.

Step 4 Select one of the following severity levels to assign to the alarm:

- Fatal
- Warning
- Critical
- Error
- Information

Step 5 Click **Create**.

Step 6 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.

Configuring RMON Controls

To change the default controls for RMON alarms, choose **Threshold Manager** from the Device Manager menu. You see the Threshold Manager window.

Click More on the Threshold Manager window. You see the second Threshold Manager dialog box.

Managing RMON Alarms

To view the alarms that have already been enabled, follow these steps:

Step 1 Choose **Events > Threshold Manager**, and then click **More** in the Threshold Manager dialog box.

Step 2 Click the **Alarms** tab. You see the RMON Alarms dialog box.

Step 3 To create a customized threshold entry, click **Create**. You see the Create RMON Alarms dialog box.

Managing RMON Event Severity Levels

To define customized RMON event severity levels, follow these steps:

Step 1 Choose **Event s> Threshold Manager**, and then click **More** in the Threshold Manager dialog box.

Step 2 Click the **Events** tab on the RMON Thresholds dialog box. You see the RMON Events dialog box.

Step 3 To create a new threshold entry, click **Create**. You see the Create Threshold Entry dialog box.

Step 4 Configure the RMON event threshold attributes.

Viewing the RMON Log

To view the RMON log from the Device Manager, follow these steps:

Step 1 Choose **Events > Threshold Manager**, and then click **More** in the Threshold Manager dialog box.

Step 2 Click the **Log** tab on the RMON Thresholds dialog box. You see the RMON Log dialog box.

Discovering SCSI Targets

This chapter describes the SCSI LUN discovery feature provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

This chapter contains the following topics:

- About SCSI LUN Discovery, page 29-1
- Initiating Customized Discovery, page 29-2
- Authenticating iSCSI Targets, page 29-2
- Specifying Targets, page 29-2
- Specifying LUN Mappings, page 29-3
- Viewing iSCSI Statistics, page 29-3
- Viewing iSCSI Sessions, page 29-3
- Viewing Session Statistics, page 29-4
- Creating an iSCSI Initiator, page 29-4
- Creating an iSCSI Virtual Target, page 29-4
- Using the iSCSI Wizard, page 29-5

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.

Initiating Customized Discovery

Customized discovery consists of a list of VSAN and domain pairs that are selectively configured to initiate a discovery. Use the **custom-list** option to initiate this discovery. The domain ID is a number from 0 to 255 in decimal or a number from 0x0 to 0xFF in hex.

Authenticating iSCSI Targets

To authenticate iSCSI targets, first specify the initiators. To specify initiators, follow these steps:

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Open Device Manager.

Step 3 Choose **IP > iSCSI**. You see the iSCSI dialog box.

Step 4 Click the **Initiators** tab if it is not already selected. You see a table listing iSCSI initiators, VSAN membership, and, if applicable, persistent node and port WWN addresses.

Step 5 Click **Create**. You see the Create iSCSI Initiators dialog box.

Step 6 Enter the initiator name in the Name field.

Step 7 Enter the VSAN membership number in the VSAN Membership field.

Step 8 Enter all the node and port information.

Step 9 Click **Create** to add this initiator to the table. Click **Close** to exit the Create iSCSI Initiators dialog box without adding the initiator. Like physical N ports, iSCSI Initiators will appear in the Fabric Login table.

Specifying Targets

To specify targets, follow these steps:

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Open Device Manager.

Step 3 Choose **IP > iSCSI**. You see the iSCSI dialog box.

Step 4 Click the **Targets** tab if it is not already selected.

You see a list of statically assigned and dynamically discovered Fibre Channel targets. Click **Import** to automatically discover and populate this list with existing targets. Click **Create** to assign a port address or control iSCSI access to certain targets.

Step 5 Click **Create**. You see the Create iSCSI Targets dialog box.

Step 6 Enter the target name in the Name field.

Step 7 Enter the port WWN, node access information, and advertised interfaces information in the appropriate fields.

Step 8 Click **Create** to add this target to the table. Click **Close** to close the Create iSCSI Targets dialog box without adding the target.

Specifying LUN Mappings

To specify LUN mappings, follow these steps:

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Open Device Manager.

Step 3 Choose **IP > iSCSI**. You see the SCSI dialog box.

Step 4 Click the **Targets** tab if it is not already selected.

Step 5 Click **Create**. You see the Create iSCSI Targets dialog box. Use this dialog box to map Fibre Channel LUNs to iSCSI LUNs:

Step 6 Enter the iSCSI LUN name in the Name field.

Step 7 Enter the iSCSI LUN, Port WWN, and FC LUN information in the appropriate fields.

Step 8 Click **Create** to add this LUN to the table. Click **Close** to close the Create iSCSI LUN Mappings dialog box without adding the LUN.

Viewing iSCSI Statistics

To view iSCSI statistics, follow these steps:

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Open Device Manager.

Step 3 Choose **IP > Statistics**. You see the Statistics dialog box.

Step 4 Click the **iSCSI** tab if it is not already selected.

Viewing iSCSI Sessions

To view iSCSI sessions, follow these steps:

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Open Device Manager.

Step 3 Choose **IP > iSCSI**. You see the iSCSI dialog box.

Step 4 Click the **Session Initiators** tab if it is not already selected.

Viewing Session Statistics

To view session statistics, follow these steps:

Step 1 Be sure you are connected to a switch that contains an IPS module.

Step 2 Open Device Manager.

Step 3 Choose **IP > iSCSI**. You see the iSCSI dialog.

Step 4 Click the **Session Statistics** tab if it is not already selected.

Creating an iSCSI Initiator

To create an iSCSI Initiator using Device Manager, follow these steps:

Step 1 Choose **IP > iSCSI**.

Step 2 Click the **Initiators** tab.

Step 3 Click **Create**. You see the Create Initiators dialog box.

Step 4 Enter the IP address, or the IQN name created from the iSCSI driver running on the initiator. The IQN name must be at least 16 characters.

Step 5 Assign names for the node WWN and port WWN fields.

There are three options. The **Auto** option assigns the WWN from a pool of about 440,000 WWNs per switch and is returned to pool when you log out. The **Persistent** option also assigns the WWN from a pool. However, when you log out of the switch, the WWIN is not returned to the pool but is saved for the initiator. The third option is to statically assign the WWN by manually entering WWN that the initiator will use.

Step 6 Click **Create** to create the initiator.

Creating an iSCSI Virtual Target

To create an iSCSI Virtual Target using Device Manager, follow these steps:

Step 1 Choose **IP > iSCSI**.

Step 2 Click the **Targets** tab.

Step 3 Click **Create**. You see the Create Targets dialog box.

Step 4 Enter the logical name to give to this virtual target.

Step 5 Click the drop-down arrow to the right of the pWWN field.

Step 6 Choose the pWWN of the FC target that will be advertised as an iSCSI virtual target.

The drop-down list shows all pWWNs that are logged into the name server.

Step 7 Choose the iSCSI initiators that will access this virtual target.

Choose All if you want all the initiators to access the target. Choose None, and then enter in the numbers (separated by commas) if you want only certain initiators to access the target.

Step 8 Click **Create** to create the virtual target.

Using the iSCSI Wizard

You can use the iSCSI Wizard to configure a new or existing initiator on a switch.

To use the iSCSI Wizard, follow these steps:

Step 1 Double-click on the **iSCSI Wizard** icon on the Fabric Manager toolbar (see Figure 29-1).

Figure 29-1 iSCSI Wizard Icon

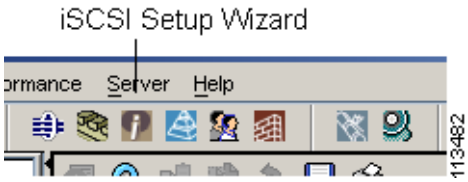

Step 2 Choose an existing initiator from the list, or enter the iqn name or IP address of the initiator, and enter the IP address of the switch on which you want to configure the new initiator.

Step 3 Assign names for the node WWN and port WWN fields.

There are three options. The **Auto** option assigns the WWN from a pool of about 440,000 WWNs per switch and is returned to pool when you log out. The **Persistent** option also assigns the WWN from a pool. However, when you log out of the switch, the WWIN is not returned to the pool but is saved for the initiator. The third option is to statically assign the WWN by manually entering WWN that the initiator will use.

Step 4 Click **Finish** to create the initiator.

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:

This chapter contains the following topics:

- About SPAN, page 30-1
- SPAN Sources, page 30-2
- SPAN Sessions, page 30-5
- Creating SPAN Sessions, page 30-5
- Editing SPAN Sources, page 30-5
- Deleting SPAN Sessions, page 30-6
- Specifying Filters, page 30-6
- SD Port Characteristics, page 30-7
- Monitoring Traffic Using Fibre Channel Analyzers, page 30-7
- Default SPAN Settings, page 30-10
- Remote SPAN, page 30-10

About SPAN

The switched port analyzer (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 non-intrusive and does not affect switching of network traffic for any SPAN source ports.

Figure 30-1 SPAN Transmission

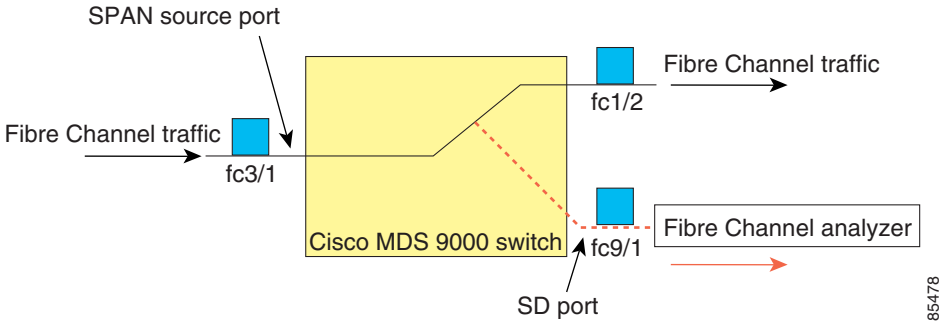

SPAN Sources

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.

Figure 30-2 SPAN Traffic from the Ingress Direction

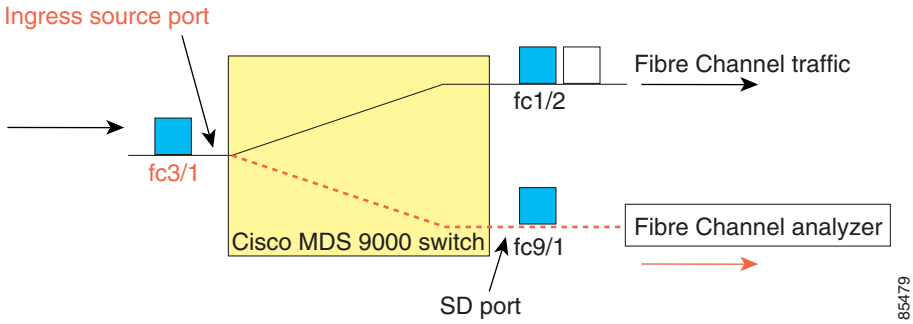

- Egress source (tx)—Traffic exiting the switch fabric through this source interface is spanned or copied to the SD port.

Figure 30-3 SPAN Traffic from Egress Direction

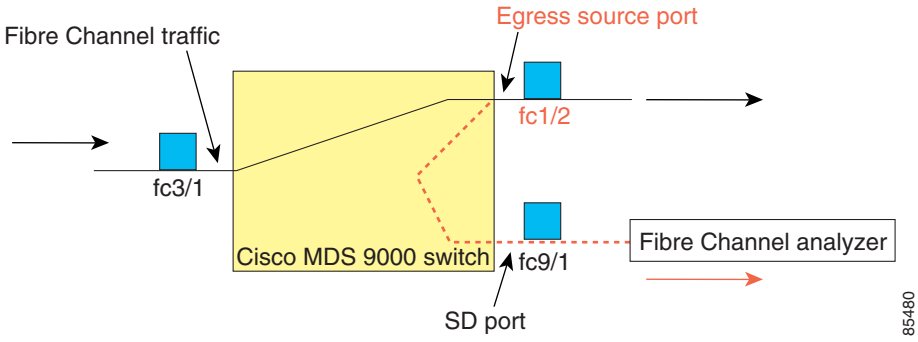

IPS Source Ports

Effective SAN-OS Release 1.3(x) Switched Port Analyzer (SPAN) capabilities are also 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 SPAN ingress traffic, egress traffic, or traffic in both directions for all eight iSCSI and 24 FCIP interfaces that are available in the IPS module.

You can configure SPAN for Ethernet traffic using Cisco switches or routers connected to the Cisco MDS 9000 Family IPS modules.

CSM Source Ports

Effective SAN-OS Release 1.3(x) Switched Port Analyzer (SPAN) capabilities are also available on the Caching Services Module (CSM).

Refer to the Cisco MDS 9000 Family SAN Volume Controller Configuration Guide for further information.

Allowed Source Interface Types

The SPAN feature is available for the following interface types:

- Physical ports
  - 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 interface
- FCIP interfaces

VSAN as a SPAN 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.
- When a VSAN is specified as a source, you will not be able to 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 SPAN source, you will not be able to 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. displays a configuration using VSAN 2 as a SPAN 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 30-4 VSAN As a SPAN Source

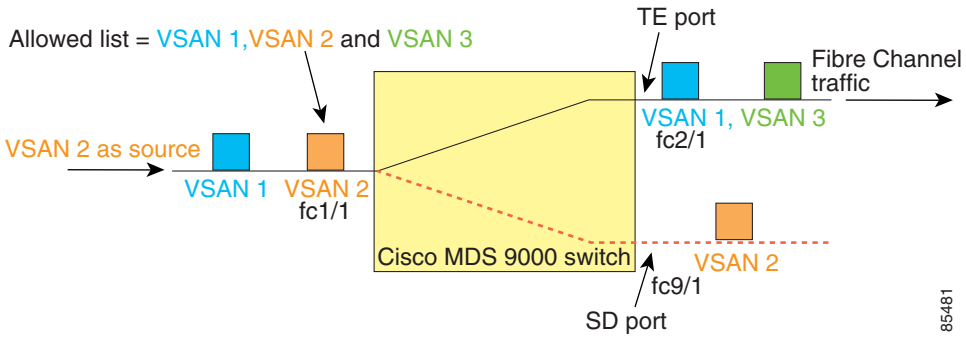

For the configuration shown in Figure 30-4, the following apply:

- VSAN 2 as a SPAN 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.

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 a SPAN session, at least one source and the SD port must be up and functioning. Otherwise, traffic will not be directed to the SD port.

To temporarily deactivate (suspend) a SPAN session use the **suspend** command in the SPAN submode. The traffic monitoring is stopped during this time. You can reactivate the SPAN session using the **no suspend** command.

A source can be shared by two sessions, however, each session must be in a different direction—one ingress and one egress.

Creating SPAN Sessions

To create a SPAN session, follow these steps.

Step 1 From the Device Manager, 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**.

Step 5 Repeat Step 4 for each session you want to create.

Step 6 Click **Close** to close the Create SPAN Session dialog box.

Step 7 Choose the destination interface by clicking once in the Dest Interface field for the appropriate session.

Step 8 Choose the filter VSAN list by clicking once in the Filter VSAN List field for the appropriate session.

Step 9 Choose active or inactive admin status by clicking the Admin drop-down menu and choosing the appropriate status.

Step 10 Click **Apply** to save your changes, or click **Close** to close the SPAN Sessions dialog box without saving your changes.

Editing SPAN Sources

To edit a SPAN source, follow these steps.

Step 1 From the Device Manager, choose **Interface > SPAN**. You see the SPAN dialog box.

Step 2 Click the **Sources** tab.

Step 3 Click once on the VSAN List field, and enter the VSAN list name.

Step 4 Click **Edit FC Source**.

You see the Edit FC Interface Source 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. Select a port and click **OK**.

Step 7 Click the direction (receive or transmit) you want.

Step 8 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 9 Click **Close**. The new FC interface source is shown in the FC Interface Source dialog box list.

Deleting SPAN Sessions

To delete a SPAN session, perform the following steps.

Step 1 From the Device Manager, choose **Interface > SPAN**.

You see the SPAN dialog box.

Step 2 Click the **Sessions** tab.

Step 3 Click once to select the SPAN session you want to delete.

Step 4 Click **Delete**.

The SPAN session is deleted.

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. Only VSANs present in the filter are spanned.

You can specify session VSAN filters which 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 buffer-to-buffer 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 port mode can not be changed if it is being used for a SPAN session.

  If you need to change a SD-port mode to another port mode, first remove the SD port from all sessions and then change the port mode using the **switchport mode** command.
- 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 Advanced Services Modules (ASMs).

Guidelines to Configure SPAN

The following guidelines apply for a SPAN configuration:

- 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.

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 when 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 30-5.

Figure 30-5 Fibre Channel Analyzer Usage Without SPAN

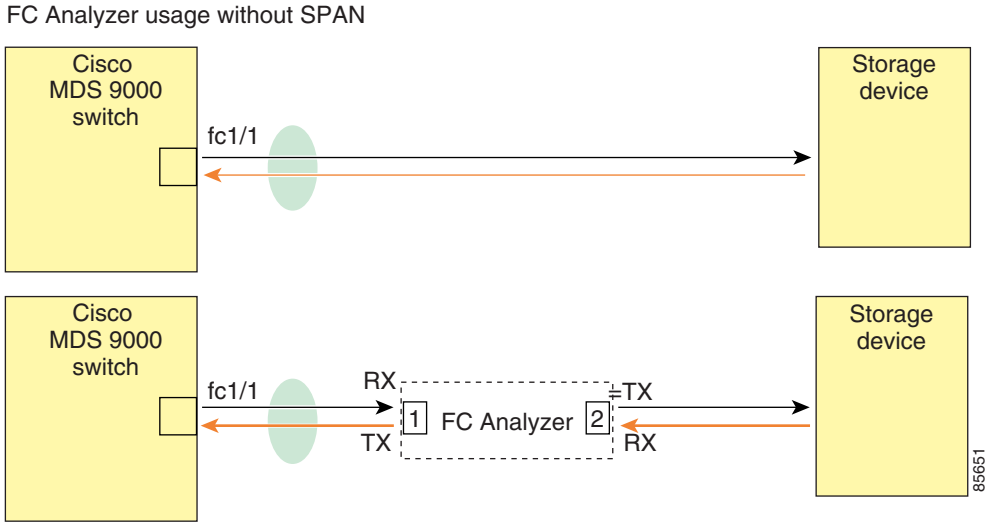

This type of connection has the following limitations:

- 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.

Using SPAN

Using SPAN, you can capture the same traffic scenario shown in Figure 30-5 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 30-6.

Figure 30-6 Fibre Channel Analyzer Using SPAN

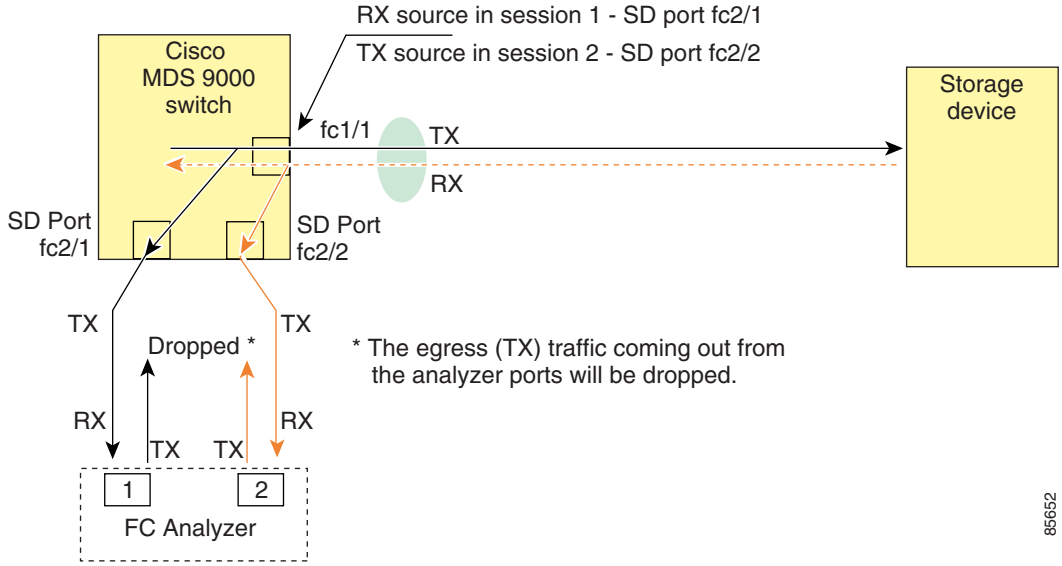

Configuring Analyzers Using SPAN

To configure Fibre Channel Analyzers using SPAN for the example in Figure 30-6, 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.

Using a 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 30-6. 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 30-7shows 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 direction. This setup is more advantageous and cost-effective than the setup shown in Figure 30-6 because it uses one SD port and one port on the analyzer, instead of using a full, two-port analyzer.

Figure 30-7 Fibre Channel Analyzer Using a Single SD Port

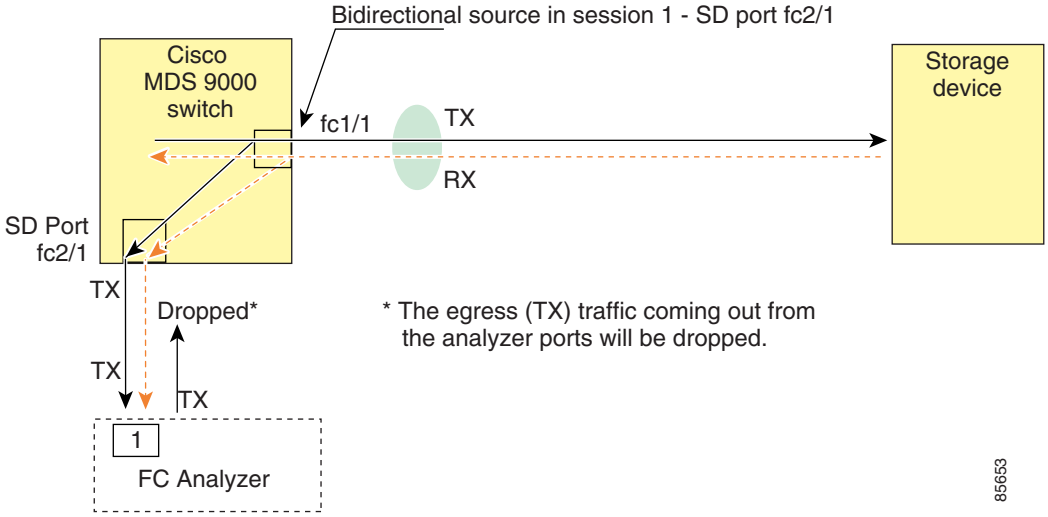

To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

Default SPAN Settings

Table 30-1 lists the default settings for SPAN parameters

Table 30-1 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.

Remote SPAN

The Remote SPAN (RSPAN) feature enables you to remotely monitor traffic for one or more SPAN sources distributed in one or more source switches in a Fibre Channel fabric. The SPAN destination (SD) port is used for remote monitoring in a destination switch. A destination switch is usually different from the source switch(es) but is attached to the same Fibre Channel fabric. You can replicate and monitor traffic in any remote Cisco MDS 9000 Family switch or director, just as you would monitor traffic in a MDS source switch.

The RSPAN feature is nonintrusive and does not affect network traffic switching for any SPAN source ports. Traffic captured on the remote switch is tunneled across a Fibre Channel fabric which has trunking enabled on all switches in the path from the source switch to the destination switch. The Fibre Channel tunnel is structured using trunked ISL (TE) ports. In addition to TE ports, the RSPAN feature uses two other interface types:

- SD ports—A passive port from which remote SPAN traffic can be obtained by the FC analyzer.
- ST ports—A SPAN tunnel (ST) port is an entry point port in the source switch for the RSPAN Fibre Channel tunnel. ST ports are special RSPAN ports and cannot be used for normal Fibre Channel traffic.

Figure 30-8 RSPAN Transmission

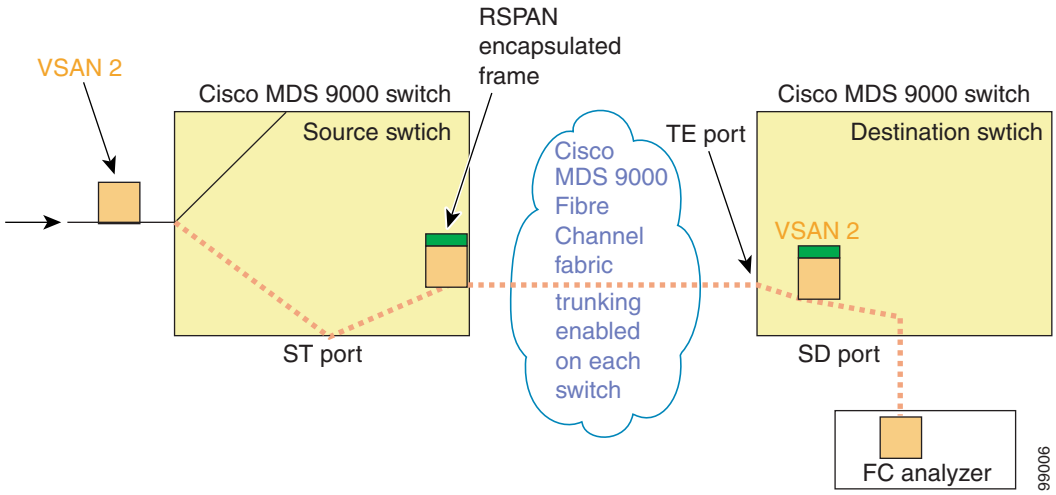

Advantages to Using RSPAN

The RSPAN features has the following advantages:

- Enables nondisruptive traffic monitoring at a remote location.
- Provides a cost-effective solution by using one SD port to monitor remote traffic on multiple switches.
- Works with any Fibre Channel analyzer.
- Compatible with the Cisco MDS 9000 Port Analyzer adapters.
- Does not affect traffic in the source switch, but shares the ISL bandwidth with other ports in the fabric.

FC and RSPAN Tunnels

A FC tunnel is a logical data path between a source switch and a destination switch. The FC tunnel originates from the source switch and terminates at the remotely located destination switch.

RSPAN uses a special Fibre Channel tunnel (FC tunnel) that originates at the ST port in the source switch and terminates at the SD port in the destination switch. You must bind the FC tunnel to a ST port in the source switch and map the same FC tunnel to a SD port in the destination switch. Once the mapping and binding is configured, the FC tunnel is referred to as a RSPAN tunnel.

Figure 30-9 FC and RSPAN Tunnel

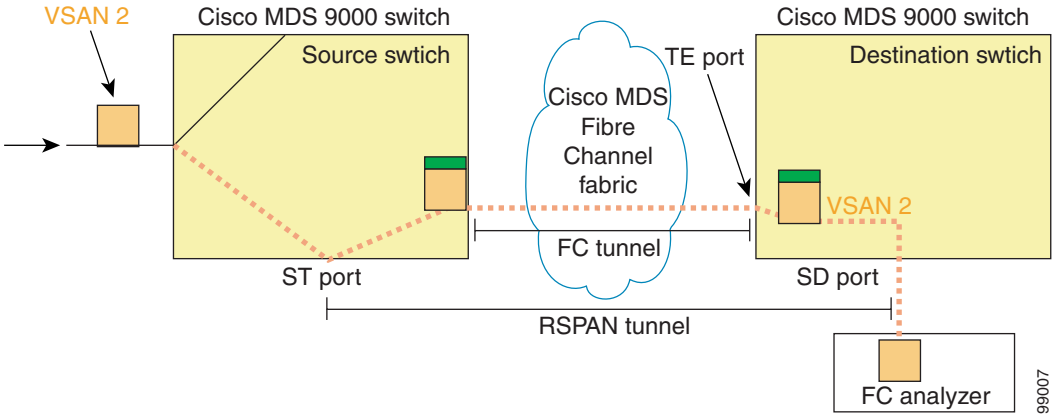

Guidelines to Configure RSPAN

The following guidelines apply for a SPAN configuration:

- All switches in the end-to-end path of the RSPAN tunnel must belong to the Cisco MDS 9000 Family.
- All VSANs with RSPAN traffic must be enabled. If a VSAN containing RSPAN traffic is not enabled, it will be dropped.
- The FC tunnel IP address must reside in the same subnet as the VSAN interface.
- The following configurations must be performed on each switch in the end-to-end path of the Fibre Channel tunnel in which RSPAN is to be implemented
  - Trunking must be enabled (the **trunk protocol enable** command is enabled by default).
  - VSAN interface must be configured (the **interface vsan** command).
  - The Fibre Channel tunnel feature must be enabled (the **fc-tunnel enable** command is disabled by default).
  - IP routing must be enabled (the **ip routing** command is disabled by default).
- If the IP address is in the same subnet as the VSAN, the VSAN interface does not have to be configured for all VSANs on which the traffic is spanned.
- A single Fibre Channel switch port must be dedicated for the ST port functionality.
- Do not configure the port to be monitored as the ST port.

ST Port Characteristics

ST port have the following characteristics:

- ST ports perform the RSPAN encapsulation of the FC frame.
- ST ports do not use BB_credits.
- An ST port can only be bound to one FC tunnel.
- ST ports cannot be used for any other purpose other than to carry RSPAN traffic.

- ST Ports cannot be configured using Advanced Services Modules (ASMs).

Configuring RSPAN

The RSPAN tunnel begins in the source switch and terminates in the destination switch. This section assumes Switch S to be the source and Switch D to be the destination.

Besides the source and destination switches, the VSAN must also be configured in each MDS switch in the Fibre Channel fabric, if they exist.

To monitor network traffic using the RSPAN feature, follow these steps:

Step 1 Create VSAN interfaces in destination switch (Switch D) and source switch (Switch S) to facilitate the Fibre Channel tunnel (FC tunnel) creation.

Step 2 Enable the FC tunnel in each switch in the end-to-end path of the tunnel.

Step 3 Initiate the FC tunnel (in Switch S) and map the tunnel to the VSAN interface IP address (in Switch D) so all RSPAN traffic from the tunnel is directed to the SD port.

Step 4 Configure SD ports for SPAN monitoring in the destination switch (Switch D).

Step 5 Configure the ST port in the source switch (Switch S) and bind the ST port to the FC tunnel.

Step 6 Create a RSPAN session in the source switch (in Switch S) to monitor network traffic.

Configuration in the Source Switch

This section identifies the tasks that must be performed in the source switch (Switch D).

This section contains the following topics:

- Creating VSAN Interfaces, page 30-13
- Configuring the ST Port, page 30-14

Creating VSAN Interfaces

Figure 30-10 depicts a basic FC tunnel configuration.

Figure 30-10 FC Tunnel Configuration

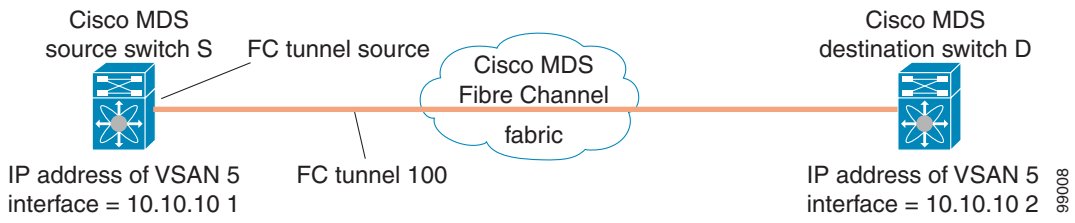

This example assumes that VSAN 5 is already configured in the VSAN database.

Configuring the ST Port

Once the FC tunnel is created, be sure to configure the ST port to bind it to the FC tunnel at the source switch. The FC tunnel becomes a RSPAN tunnel once the binding and mapping is complete. depicts a basic FC tunnel configuration.

Figure 30-11 Binding the FC Tunnel

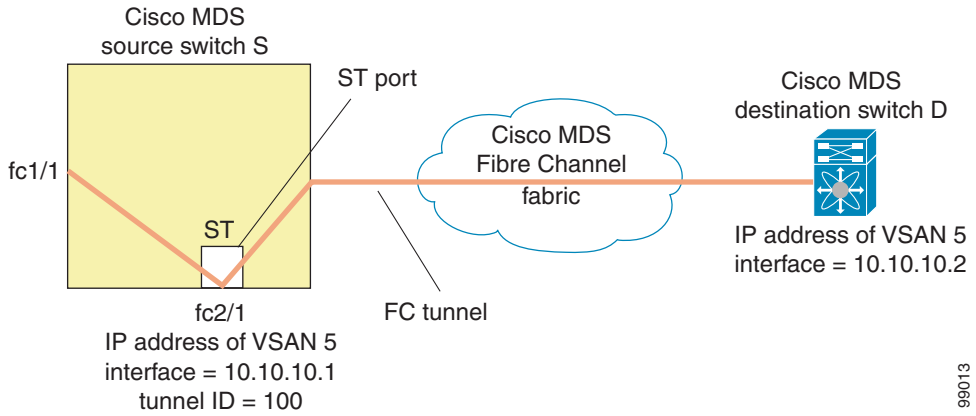

ST ports cannot be configured using Advanced Services Modules (ASMs).

Configuration in All Intermediate Switches

This section identifies the tasks that must be performed in all intermediate switches in the end-to-end path of the RSPAN tunnel.

This section contains the following topics:

- Configuring VSAN Interfaces, page 30-14
- Enabling IP Routing, page 30-14

Configuring VSAN Interfaces

Figure 30-12 depicts an RSPAN tunnel configuration terminating in the destination switch (Switch D).

This example assumes that VSAN 5 is already configured in the VSAN database.

Enabling IP Routing

The IP routing feature is disabled by default. Be sure to enable IP routing in each switch (including the source and destination switches) in the end-to-end path in the fabric. This step is required to setup the FC tunnel.

Configuration in the Destination Switch

This section identifies the tasks that must be performed in the destination switch (Switch D).

This section contains the following topics:

- Configuring the SD Port, page 30-15
- Mapping the FC Tunnel, page 30-15

Configuring the SD Port

The SD port in the destination switch enables the FC Analyzer to receive the RSPAN traffic from the Fibre Channel tunnel. depicts a RSPAN tunnel configuration, now that tunnel destination is also configured.

Figure 30-12 RSPAN Tunnel Configuration

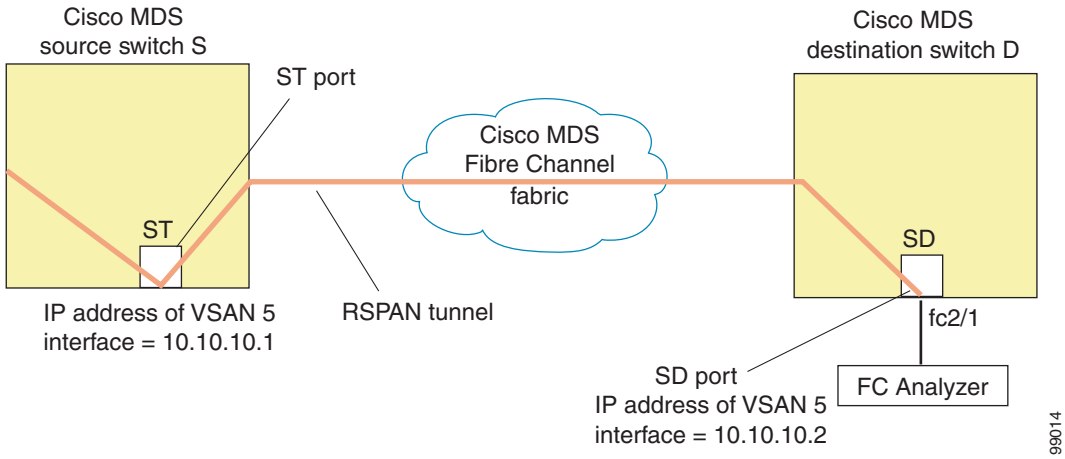

SD ports cannot be configured using Advanced Services Modules (ASMs).

Mapping the FC Tunnel

The **tunnel-id-map** option specifies the egress interface of the tunnel at the destination switch.

Figure 30-13 FC Tunnel Configuration

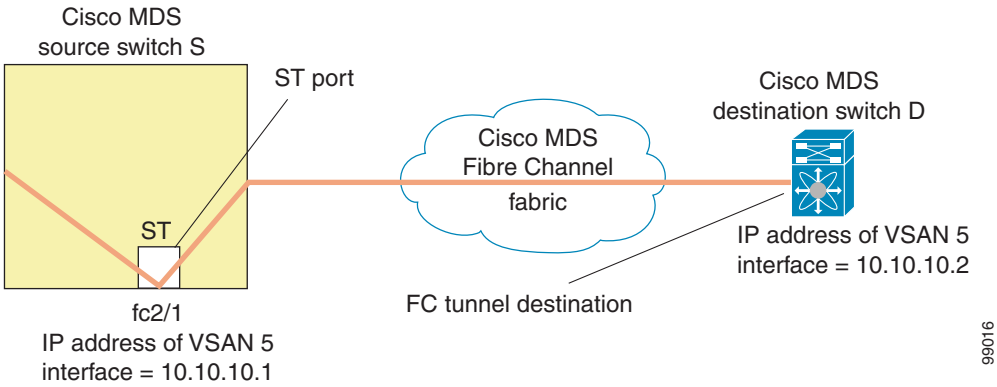

Configuring An Explicit Path

You can specify an explicit path through the Cisco MDS Fibre channel fabric (source-based routing), use the **explicit-path** option. For example, if you have multiple paths to a tunnel destination, you can use this option to specify the fc-tunnel to always take one path to the destination switch. The software then use this specified path even if other paths are available.

This option is especially useful if you prefer to direct the traffic through a certain path although other paths available. In a RSPAN situation, you can specify the explicit-path so the RSPAN traffic does not interfere with the existing user traffic. You can create any number of explicit paths in a switch.

Figure 30-14 Explicit Path Configuration

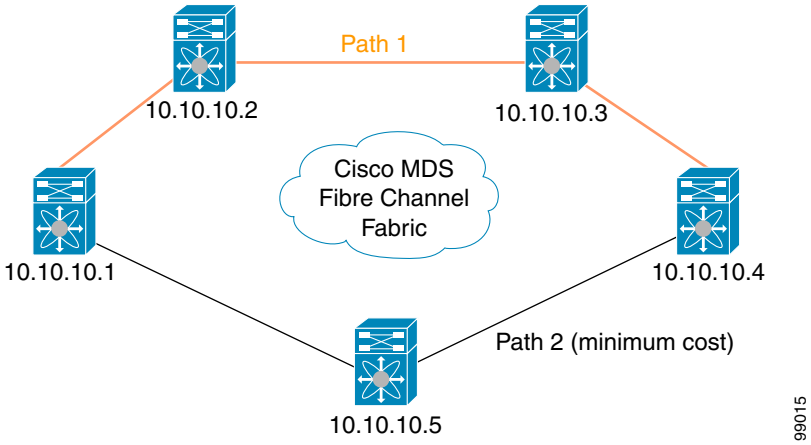

The explicit path must be created in the source switch. To configure an explicit path, you must first create the path and then configure the use of any one path. If an explicit path is not configured, the minimum cost path is used by default. If an explicit path is configured and is functioning, the specified path is used.

This configuration explicitly specifies Path 1 to be used for the RSPAN traffic. Refer to RFC 3209 for further details on explicit paths and source based routing.

Monitoring RSPAN Traffic

Once the session is configured, other SPAN sources for this session can also be configured as required. shows a RSPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress direction.

Figure 30-15 Fibre Channel Analyzer Using a Single SD Port to Monitor RSPAN Traffic

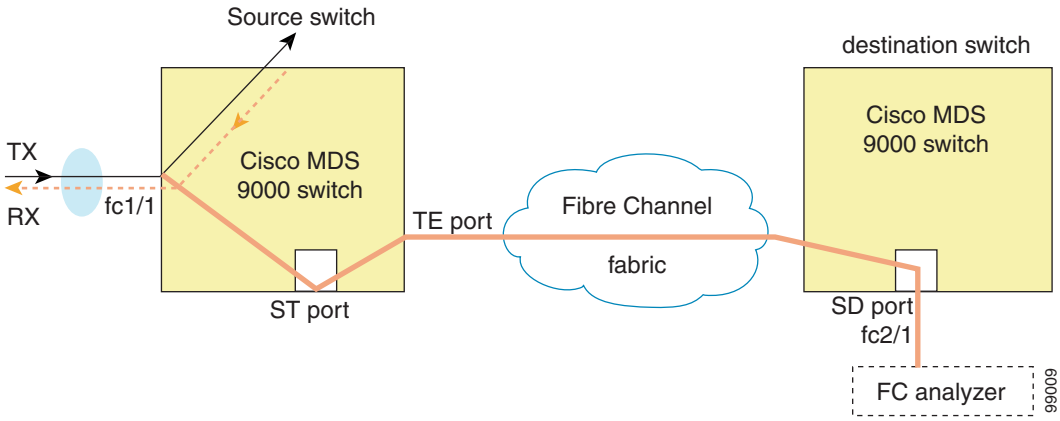

To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

Sample Scenarios

RSPAN can be combined with the local SPAN feature so SD ports forward local SPAN traffic along with remote SPAN traffic. Various SPAN source and tunnel scenarios are described in this section.

Single Source with One RSPAN Tunnel

The source Switch S and the destination Switch D are interconnected through a Fibre Channel fabric. A RSPAN tunnel is configured as a destination interface for SPAN session and the ST port forwards SPAN traffic through the RSPAN tunnel.

Figure 30-16 RSPAN Scenario with One Source Switch, One Destination Switch, and One Tunnel

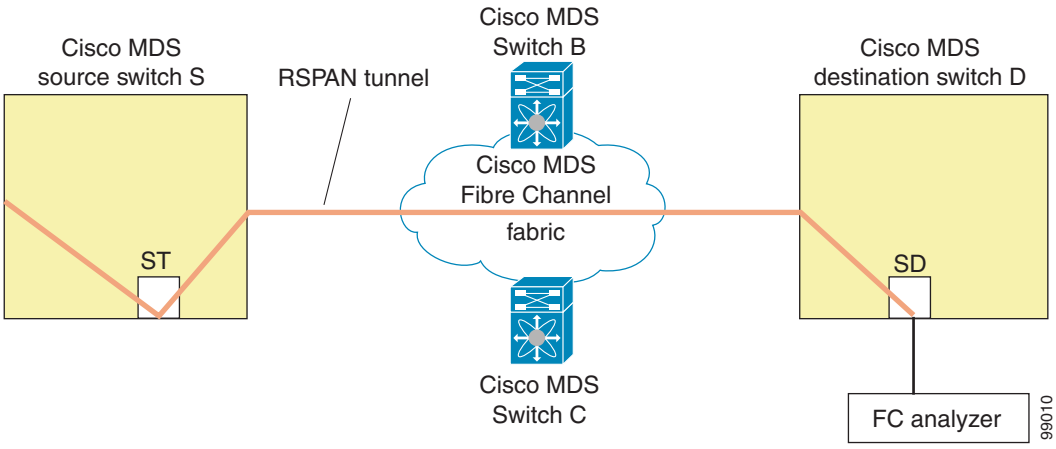

Single Source with Multiple RSPAN Tunnels

Figure 30-17 displays two separate RSPAN tunnels configured between Switches S and D. Each tunnel has an associated ST port in the source switch and a separate SD port in the destination switch. This configuration is useful for trouble shooting purposes.

Figure 30-17 RSPAN Scenario with One Source Switch, One Destination Switch, and Multiple Tunnels

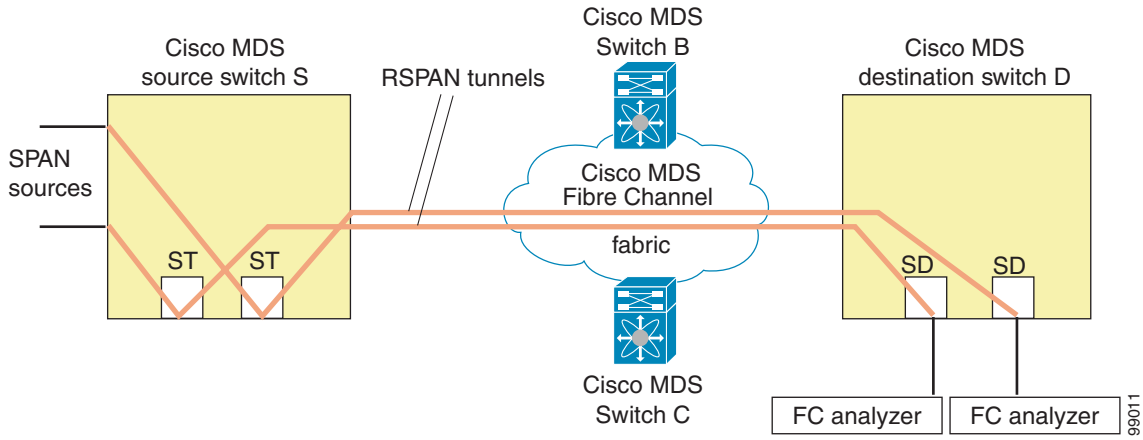

Multiple Sources with Multiple RSPAN Tunnels

Figure 30-18 displays two separate RSPAN tunnels configured between Switches S1 and S2. Both tunnels have an associated ST port in their respective source switch and terminate in the same SD port in the destination switch.

Figure 30-18 RSPAN Scenario with Two Source Switches, a Destination Switch, and Multiple Tunnels

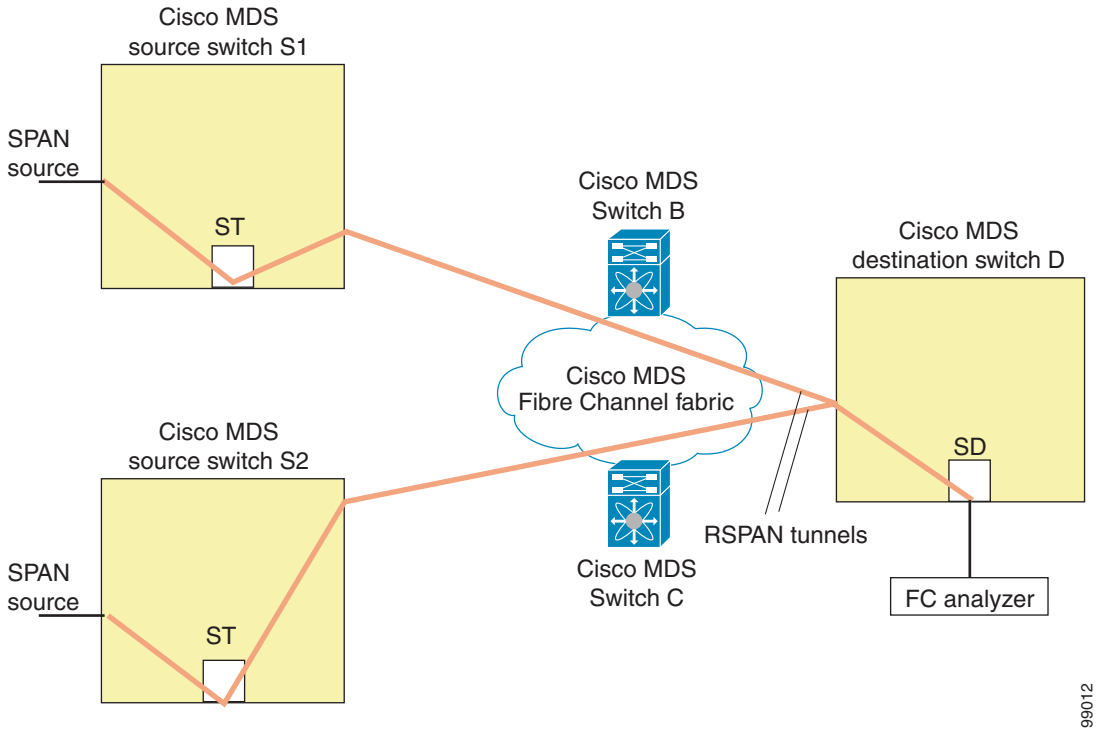

This configuration is useful for remote monitoring purposes. For example, the administrator may be at the destination switch and can remotely monitor the two source switches.

Advanced Features and Concepts

This chapter describes the advanced features provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

This chapter contains the following topics:

- Configuring FC Timers, page 31-1
- Configuring a Fabric Analyzer, page 31-2
- Configuring World Wide Names, page 31-5
- Allocating Flat FC IDs, page 31-5
- Enabling Loop Monitoring, page 31-6
- Configuring the Switch for Interoperability, page 31-6
- Using the show tech-support Command, page 31-8
- Managing World Wide Names, page 31-9
- Configuring Timers, page 31-9

Configuring FC Timers

The **fctimer** command modifies Fibre Channel protocol related timer values for the switch.

You can use the **fctimer** command in configuration mode to configure the following 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.

Configuring Timers Per-VSAN

You can also issue the **fctimer** command 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 or activated when their timer values are changed.

Caution You cannot perform a nondisruptive downgrade to any earlier version which 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 SAN-OS Release 1.2(x) or 1.1(x) after the timer is configured for a VSAN, an error message is issued to warn the user about strict incompatibilities. Refer to the Cisco MDS 9000 Family Troubleshooting Guide for further information.

Configuring a Fabric Analyzer

Fibre Channel protocol analyzers capture, decode, and analyze frames and ordered sets on a link. While 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 capability 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—You can obtain more information from http://www.tcpdump.org.
- Ethereal—You can obtain more information from 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.

This section explains the following topics:

This section contains the following topics:

- About the Cisco Fabric Analyzer, page 31-3
- Configuring the Cisco Fabric Analyzer, page 31-4

About the Cisco Fabric Analyzer

The Cisco Fabric Analyzer comprises two separate components:

- A 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
- A 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 31-1 Cisco Fabric Analyzer Usage

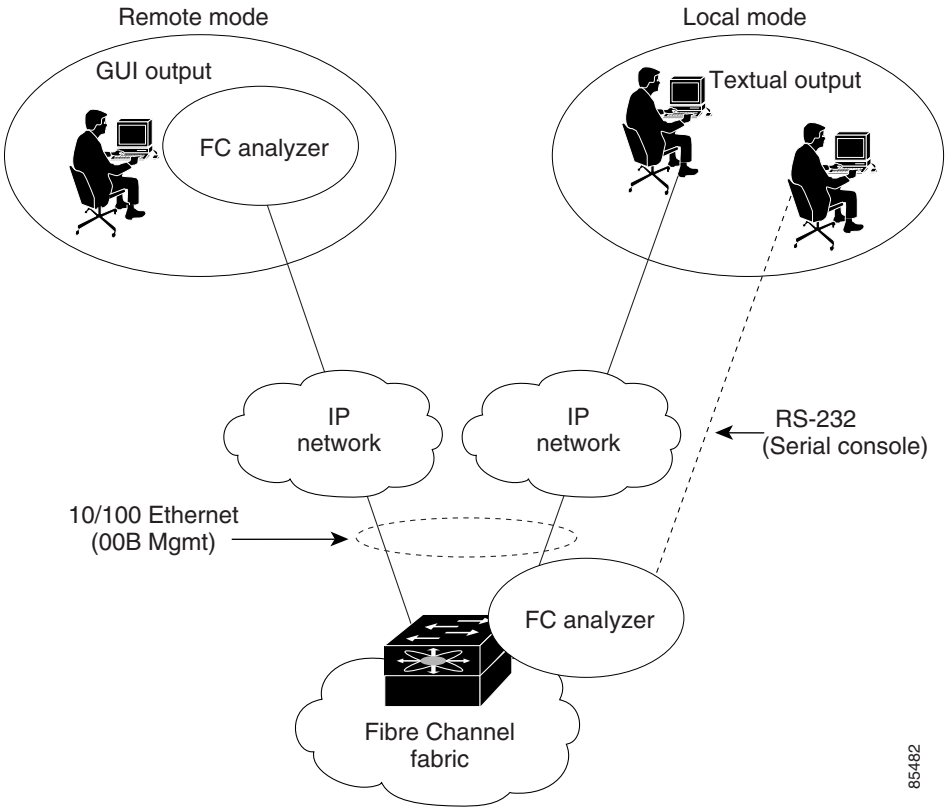

Local Text-Based Capture

This component is a command-line driven text-based interface that captures traffic to and from the supervisor module in a Cisco MDS 9000 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 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 end points, 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 unauthorized machine in the network from snooping on the control traffic in the network.

RPCAP supports two setup connection modes based on fire wall 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. Since Ethereal has a GUI front-end, it supports a rich functionality such as colorized display, graphical assists in defining filters, and searching for specific frames. These features are documented on the Ethereal web site.

While remote capture via 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 by issuing the **fcanalyzer local** or **fcanalyzer remote** commands in configuration mode.

- Local capture—The command setting to enable a local capture cannot be saved to persistent storage or synchronized to standby.
- Remote capture—The command setting to enable a remote capture can be saved to persistent storage using the **copy** command. 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.

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 supervisor module, assigns WWNs to each switch. This WWN is independent of other WWNs on each switch. This centralized control of WWN has the following advantages:

- Efficient sharing of WWN space
- Centralized support across switches

Cisco MDS 9000 Family switches support three network address authority (NAA) address formats.

Table 31-1 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 worldwide names should be made by an administrator or individual who is completely familiar with switch operations.

Allocating Flat FC IDs

Fibre Channel standards require a unique FC ID to be allocated to a 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.

Based on the assigned FC ID, some HBAs assume that no other ports have the same area bits and domain. When a target is assigned with a FC ID that has the same area bits, but different port bits, the HBA fails to discover these targets. To isolate these HBAs in a separate area, switches in the Cisco MDS 9000 Family follow a different FC ID allocation scheme. By default, the FC ID allocation mode is **auto**. In the **auto** mode, only HBAs without interop issues are assigned FCIDs with specific ports bits. All other HBAs are assigned FC IDs with a whole area (port bits set to 0).

The three options to allocate FCID are auto (default), none, and flat.

Caution Changes to FC IDs should be made by an administrator or individual who is completely familiar with switch operations.

Enabling Loop Monitoring

When a disk is removed from a loop port, the loop stays active based on the bypass circuit. Thus the disk removal is not known until you try to communicate with the disk. To detect such removals, the disks can be polled periodically (every 20 seconds) using the **fcinterop loop-monitor** command. This command enables loop polling for FL ports in a Cisco MDS 9000 Family switch. By default, the **fcinterop loop-monitor** command is disabled.

Caution Changes to the loop monitoring feature should be made by an administrator or individual who is completely familiar with switch operations.

Configuring the Switch for Interoperability

Interoperability enables multiple vendor products 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 provide the product with a more amiable standards compliant implementation.

Table 31-2 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 31-2 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 the McData nominal restriction to this same range. They can either be setup statically (the MDS will only accept 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	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 31-2 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 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 doesn't 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 using the **copy running start** command.
Zone propagation	Some vendors do not pass the full zone configuration (**zoneset**) to other switches, only the active zoneset gets passed. Verify that the active zoneset 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 Port-Channels cannot be used to connect MDS to non-MDS switches. Only E ports can be used to connect to non-MDS switches. TE ports and PortChannels can still be used to connect an MDS to other 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	Prior to Cisco MDS SAN-OS Release 1.3(4a), IVR-enabled VSANs can only be configured in no interop (default) mode or in interop mode 1. As of Cisco MDS SAN-OS Release 1.3(4a), IVR-enabled VSANs can be configured in no interop (default) mode or in any interop mode.

Configuring Interoperability

The **interop** mode in Cisco MDS 9000 Family switches can be enabled disruptively or nondisruptively.

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 and McData switches do not understand. Rejecting these frames causes the common E ports to become isolated.

Using 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 of this command can be provided to technical support representatives when reporting a problem.

The **show tech-support** command displays the output of several **show** commands at once. The output from this command will vary depending on your configuration. Use the **show tech-support** command in EXEC mode to display general information about the switch when reporting a problem.

You can choose to have detailed information for each command or even specify the output for a particular interface, module or VSAN. Each command output is separated by line and the command precedes the output.

Note Explicitly set the **terminal length** command to 0 (zero) to disable auto-scrolling and enable manually scrolling. Use the **show terminal** command to view the configured the terminal size. After obtaining the output of this command, remember to reset you terminal length as required

Tip You can save the output of this command to a file by appending **>** filename to the **show tech-support** command If you save this file, verify you have sufficient space to do so--each of these files may take about 1.8 MB. However, you can zip this file using the **gzip** filename command Copy the zipped file to the required location using the **copy** command and unzip the file using the **gunzip** command

The default output of the **show tech-support** command includes the output of the following commands:

- **show version**
- **show environment**
- **show module**
- **show hardware**
- **show running-config**
- **show interface**
- **show accounting log**
- **show process**
- **show process log**
- **show processes log details**
- **show flash**

Each command is discussed in both the *Cisco MDS 9000 Family Configuration Guide* and the *Cisco MDS 9000 Family Command Reference*. Refer to the *Cisco MDS 9000 Family Troubleshooting Guide* to obtain debug processes, procedures, and examples.

Managing World Wide Names

Each port on a Cisco MDS 9000 Family switch is uniquely identified by its world wide names (WWNs), which include the switch MAC address and an identifier for each port. The principal switch selection and the allocation of domain IDs use the WWN to identify a specific port.

To add WWNs, perform the following steps.

Step 1 From the Fabric Manager, choose **FC > WWN Manager** on the menu tree, or from the Device Manager, choose **FC > WWN Manager**.

The Information pane of the Fabric Manager displays WWN information for multiple switches. The dialog box from the Device Manager displays WWN information for a single switch.

Step 2 Configure the BaseMacAddress and MacAddressRange attributes for the WWN(s).

Step 3 In the Fabric Manager Information pane, the information is updated. In the Device Manager dialog, click **Apply** to accept the changes; click **Close** to close the WWWN Manager dialog without saving changes.

Configuring Timers

To configure timers from the Fabric Manager, choose FC > Timers & Policies on the menu tree. The Information pane from the Fabric Manager displays timers for multiple switches.

To configure timers from the Device Manager, choose FC > Timers/Policies. The dialog box from the Device Manager displays timers for a single switch.

Configuring Fabric Configuration Servers

This chapter describes the Fabric Configuration Server (FCS) feature provided in the Cisco MDS 9000 Family of directors and switches.

This chapter contains the following topics:

- About FCS, page 32-1

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 its 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 their 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 is a 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 SNMP or CLI.

In Figure 32-1, Management Application 1 (M1) is connected through an F port with a port VSAN ID 1 and Management Application 2 (M2) is connected through an F port with a port VSAN ID 2. M1 can query FCS info 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 32-1 FCSs in a VSAN Environment

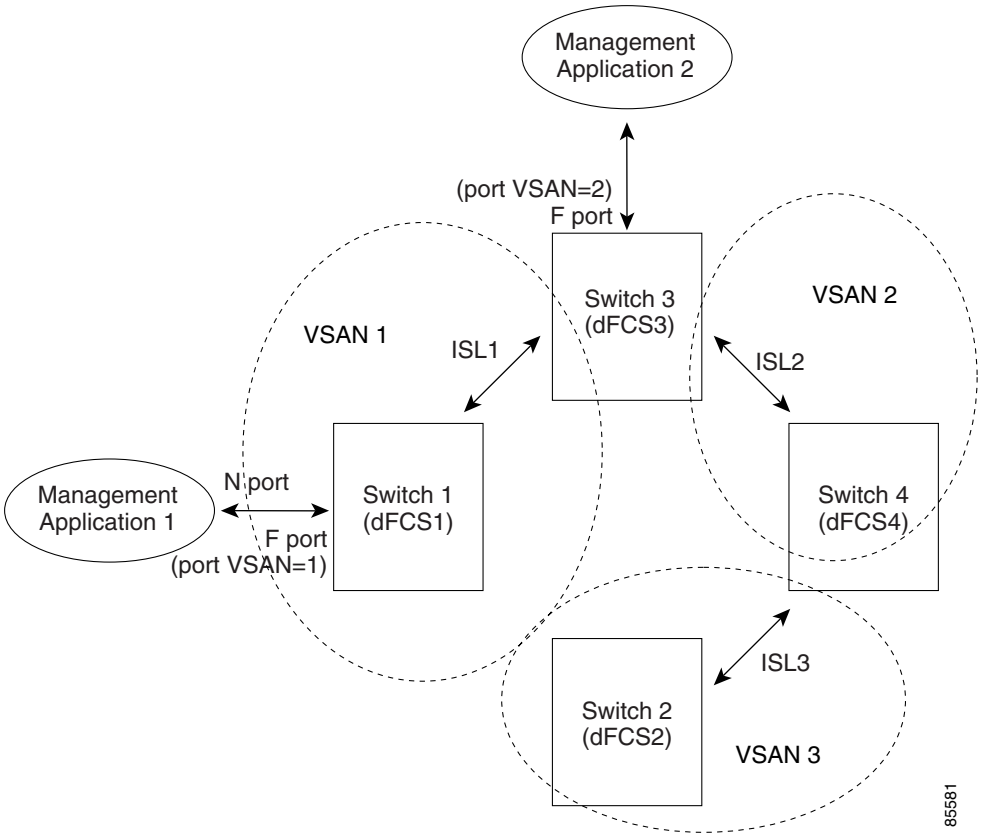

Significance of FCS

The significance of FCSs are as follows:

- Network Management
  - N port management application can query and obtain information about the fabric elements.
  - SNMP Manager--A SNMP Manager can use the FCS management information base (MIB) to start discovery and obtain information about the fabric topology.
- FCSs supports 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 updates it with every change. When a restart or switchover happens, FCSs retrieve the secondary storage information, and rebuild its database.
- The SNMP manager can query FCSs for all the IEs, ports, and platforms in the fabric.

Monitoring System Processes and Logs

This chapter provides details on monitoring the health of the switch. It includes the following sections:

This chapter contains the following topics:

- Configuring Kernel Core Dumps, page 33-1

Configuring Kernel Core Dumps

Caution Changes to the kernel cores should be made by an administrator or individual who is completely familiar with switch operations.

When a specific module operating system (OS) crashes, it is sometimes useful to obtain a full copy of the memory image (called a kernel core dump) to identify the cause of the crash. When the module experiences a kernel core dump it triggers the proxy server configured on the supervisor. The supervisor sends the module OS kernel core dump to the Cisco MDS 9000 System Debug Server. Similarly, if the supervisor OS fails the supervisor sends its OS kernel core dump to the Cisco MDS 9000 System Debug Server.

The Cisco MDS 9000 System Debug Server is a Cisco application that runs on Linux. It creates a repository for kernel core dumps. You can download the Cisco MDS 9000 System Debug Server from the Cisco.com website.

Kernel core dumps are only useful to your technical support representative. The kernel core dump file, which is a large binary file, must be transferred to an external server that resides on the same physical LAN as the switch. The core dump is subsequently interpreted by technical personnel who have access to source code and detailed memory maps.

Tip Core dumps take up disk space on the Cisco MDS 9000 System Debug Server application. If all levels of core dumps (**level all** option) are configured, you need to ensure that a minimum of 1GB of disk space is available on the Linux server running the Cisco MDS 9000 System Debug Server application to accept the dump. If the process does not have sufficient space to complete the generation, the module resets itself.

All changes made to kernel cores are saved to the running configuration

Troubleshooting the Fabric

There are several things you can do to use Fabric Manager to troubleshoot your fabric.

This chapter contains the following topics:

- Analyzing Switch Device Health, page 34-1
- Analyzing End-to-End Connectivity, page 34-2
- Analyzing Switch Fabric Configuration, page 34-2
- Analyzing the Results of Merging Zones, page 34-3
- Issuing the Show Tech Support Command, page 34-3
- Using Traceroute and Other Troubleshooting Tools, page 34-4
- Locating Other Switches, page 34-5
- Configuring an OUI, page 34-5

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, follow these steps:

Step 1 Choose **Tools > Switch Health** from the Fabric Manager. You see the Switch Health Analysis window.

Step 2 Click **Start** to identify any problems that may currently be affecting the selected switch. The Switch Health Analysis window displays any problems affecting the selected switches.

Step 3 Fix these problems.

Step 4 Click **Clear** to remove the contents of the Switch Health Analysis window.

Step 5 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.

To use this option, follow these steps:

Step 1 Choose **Tools > End to End Connectivity** from the Fabric Manager. You see the End to End Connectivity window.

Step 2 Choose the VSAN in which you want to verify connectivity from the VSAN drop-down list.

Step 3 Identify any latency issues in the network fabric by clicking the option **Report average latencies greater than** and entering the number of microseconds.

Step 4 Click **Ensure that members can communicate** to perform a Fibre Channel ping between the selected end points.

Step 5 Identify the number of packets, the size of each packet, and the timeout in milliseconds.

Step 6 Analyze the redundant paths between endpoints by clicking **Ensure that redundant paths exist between members**.

Step 7 Click **Analyze**. The End to End Connectivity Analysis window displays the selected end points with the switch to which each is attached, and the source and target ports used to connect it.

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

- Ignoring empty zone—No requests are issued for this zone.
- Ignoring zone with single member—No requests are issued for this zone.
- Source/Target are unknown—No nameserver entries exist for the ports or we have not discovered the port during discovery.
- Both devices are on the same switch.
- 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 8 Click **Clear** to remove the contents of the window.

Step 9 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 to analyze the configuration of a switch, follow these steps:

Step 1 Choose **Tools > Fabric Configuration** from the Fabric Manager. You see the Fabric Configuration window.

Step 2 Choose if you want to compare the selected switch to another switch or to a Policy File.

- If you are making a switch comparison, click **Switch** and then click the drop-down arrow to see a list of switches.
- If you are making a policy comparison, click **Policy File**, and the 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 default rules as required and click **OK**.

Step 5 Click **Compare**. The system analyzes the configuration and displays issues that arise as a result of the comparison.

Step 6 Click to place a check mark in the Resolve column for the issues you want to resolve.

Step 7 Resolve them by choosing the **Resolve Issues** option.

Step 8 Click **Clear** to remove the contents of the window.

Step 9 Click **Close** to close the window.

Analyzing the Results of Merging Zones

You can use the Zone Merge option on the Fabric Manager Tools menu to determine if two connected switches have compatible zone configurations.

To use the Zone Merge option, follow these steps:

Step 1 Choose **Tools > Zone Merge** from the Fabric Manager. You see the Zone Merge Analysis window.

Step 2 Choose a switch from each pull-down list.

Step 3 Identify the VSAN for which you want to perform the zone merge analysis.

Step 4 Click **Analyze**. The Zone Merge Analysis window displays any inconsistencies between the zone configuration of the two selected switches.

Step 5 Click **Clear** to remove the contents of the window.

Step 6 Click **Close** to close the window.

Issuing the Show Tech Support Command

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 use the Fabric Manager **show tech support** command, perform the following steps.

Step 1 Choose **Tools > Show Tech Support**. You see the Show Tech Support dialog box.

Step 2 Choose the switches for which you want to view Show Tech Support information by checking the check boxes next to their IP addresses.

Step 3 Choose the directory where you want the text files (containing the Show Tech Support information) to be written.

Step 4 Enter your username and password in the appropriate fields.

Note In order for Fabric Manager to successfully issue the **show tech support** command on a switch, that switch must have this username and password. Fabric Manager will be unable to log into a switch that does not have this username and password, and an error will be returned for that switch.

Step 5 Set the timeout value. The default is 30 seconds.

Step 6 Check the SSH check box if you want to use SSH to connect to the switch. If you do not check the SSH check box, Telnet is used. (SSH is slower than Telnet, so if you are using SSH you may want to increase the timeout value described in Step 5.)

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.

In the Status column next to each switch, a highlighted status is displayed. 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 indicates that the Show Tech Support command has completed successfully. On successful completion, a button becomes available in the View column for each switch.

Step 8 To view the Show Tech Support output, click the button next to the name of the switch. You see the Show Tech Support information in your default text editor.

Note If you would like to view the Show Tech Support files without using Fabric Manager, you can open them with any text editor. Each file is named with the switch IP address and has a .TXT extension (for example, 111.22.33.444.txt).

Using Traceroute and Other Troubleshooting Tools

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

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

To use the Traceroute option to verify connectivity, follow these steps:

Step 1 Select two or more endpoints on the Fabric Manager map.

Step 2 Choose **Tools > Traceroute**, or right-click one of the endpoints and choose **Trace Route** from the pop-up menu. You see the Traceroute window.

Step 3 Change the timeout value if the default (10 seconds) is too short or too long.

Step 4 Click **Start**.

Step 5 The results of the Traceroute operation appear in the Results box.

Locating Other Switches

The Locate Switches option uses SNMPv2 and discovers devices responding to SNMP requests with the read-only community string public.

To locate switches that are not included in the currently discovered fabric, follow these steps:

Step 1 Choose **File > Locate Switches** from the Fabric Manager main window. You see the Locate Switches dialog box.

Step 2 Enter a range of specific addresses belonging to a specific subnet which 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.

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.

Configuring an OUI

When two WWNs in different VSANs on the same fabric have the same IP address, you will need to specify an Organizationally Unique Identifier (OUI) that Fabric Manager can use to differentiate the WWNs. If Fabric Manager encounters this situation, a dialog is displayed. Enter the OUI in the appropriate fields in the dialog. Restart Fabric Manager for your changes to take effect.

Note This situation does not affect the availability or the functionality of the switch and/or fabric.

Troubleshooting Fabric Manager Issues

This chapter contains some common issues you may experience while using Cisco Fabric Manager, and provides solutions.

This chapter contains the following topics:

- Can I Set the Map Layout So It Stays After I Restart Fabric Manager?, page 35-1
- Two Switches Show on my Map, But I Only Have One Switch, page 35-2
- There is a Red Line Through the Switch. What's Wrong?, page 35-2
- There is a Dotted Orange Line Through the Switch. What's Wrong?, page 35-2
- Can I Upgrade Without Losing My Map Settings?, page 35-2
- Are There Any Restrictions When Using Fabric Manager Across FCIP?, page 35-2
- Running Cisco Fabric Manager with Multiple Interfaces, page 35-3
- Configuring a Proxy Server, page 35-4
- Clearing Topology Maps, page 35-5
- Can I Use Fabric Manager in a Mixed Software Environment?, page 35-5

Can I Set the Map Layout So It Stays After I Restart Fabric Manager?

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, do the following:

Step 1 Right-click on a blank space in the map. A menu is displayed.

Step 2 Select Layout -> Fix All Nodes from the menu.

Two Switches Show on my Map, But I Only Have 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 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 config checker will do this task.)

There is a Red Line Through the Switch. What's Wrong?

If a red line shows through your switch, this means Fabric Manger sees something wrong with the switch. Check the Switch->Inventory report. A module, fan, or power supply has failed or is offline and plugged in.

There is a Dotted Orange Line Through the Switch. What's Wrong?

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.

Can I Upgrade Without Losing My Map Settings?

When you upgrade from one version of Fabric Manager to another, there is a way to prevent the loss of map settings (enclosure names, placement on the map, etc.)

The $HOME/.cisco_mds9000/db directory contains all discovered fabrics (*.dat) and maps (*.map). These are upgradable between 1.1 and 1.2. If you need to clear the fabric cache, you should first export the enclosures to a file to avoid losing them. Everything else aside from enclosures and map coordinates are stored on the switch. The preferences, last opened, and site_ouis.txt format doesn't change from release to release.

Are There Any 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.

Running 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.

- Specifying an Interface for Fabric Manager Server, page 35-3
- Specifying an Interface for Fabric Manager Client or Device Manager, page 35-4
- Specifying an Interface for Performance Manager, page 35-3

Specifying an Interface for Fabric Manager Server

To specify an interface for Fabric Manager Server, perform the following steps:

Step 1 Go to the .cisco_mds9000 folder.

Step 2 Edit the server.properties file with a text editor.

Step 3 Scroll until you find the line snmp.localaddress

Step 4 If the line is commented, remove the comment character.

Step 5 Set this value to the IP address or interface name of the NIC you want to use.

Step 6 Save the file.

Step 7 Stop and restart Fabric Manager Server.

Specifying an Interface for Performance Manager

To specify an interface for Performance Manager, perform the following steps:

Step 1 Go to the .cisco_mds9000 folder.

Step 2 Edit the PMCollector.conf file with a text editor.

Step 3 Scroll until you find the line wrapper.java.additional.2=-Dmds.nmsAddress=

Step 4 If the line is commented, remove the comment character.

Step 5 Set this value to the IP address or interface name of the NIC you want to use.

Step 6 Save the file.

Step 7 Stop and restart Performance Server.

Specifying an Interface for Fabric Manager Client or Device Manager

To specify an interface for the Fabric Manager Client or Device Manager, perform the following steps:

Step 1 Go to the .cisco_mds9000/bin folder.

Step 2 Edit the DeviceManager.bat file or the FabricManager.bat file.

Step 3 Scroll to the line that begins with set JVMARGS=

Step 4 Add the parameter -Dmds.nmsAddress=ADDRESS, where ADDRESS is the IP address or interface name of the NIC you want to use.

Step 5 Save the file and relaunch Fabric Manager Client or Device Manager.

Configuring a Proxy Server

If your network uses a proxy server for HTTP requests, make sure the Java Web Start Application Manager is properly configured with the IP address of your proxy server.

To configure a proxy server in the Java Web Start Application Manager, follow these steps:

Step 1 Double-click the **Java Web Start application manager** icon on your Windows desktop, or choose **Program Files > Java Web Start**.

Step 2 Select **File > Preferences** from the Java WebStart Application Manager.

Step 3 Click the **Manual** radio button and enter the IP address of the proxy server in the HTTP Proxy field.

Step 4 Enter the HTTP port number used by your proxy service in the HTTP Port field.

Step 5 Click **OK**.

Clearing Topology Maps

If you have a switch that you have removed from the fabric, there will be a red X through the switch's icon. You can clear this information from the Fabric Manager client, or from the Fabric Manager server (which will clear the information for all clients) without having to reboot the switch.

To clear information from topology maps, follow these steps:

Step 1 In the Map pane, click on the **Refresh Map** icon. This clears the information from the client.

Step 2 Choose **Server > Purge**. This clears the information from the server.

Can I Use Fabric Manager in a Mixed Software Environment?

You can use Fabric Manager version 1.3(x) to manage a mixed fabric of Cisco MDS 9000 switches. Certain 1.3 feature tabs will be empty for any switches running a software version that does not support those features.

A

authentication
  global overide 24-37

D

default zones
  interoperability 31-7
document conventions xxx
domain IDs
  interoperability 31-6

E

Ethereal freeware
  analyzer 31-4
  information 31-2

F

Fabric Analyzer
  configuring 31-4
  description 31-2
FC IDs
  allocating 31-5
  allocating areas 31-5
FSPF
  interoperability 31-7

I

interoperability
  configuring 31-6
iSCSI targets
  secondary access 24-29

L

libpcap freeware 31-2
local capture 31-4
loop monitoring 31-6
loop port 31-6

N

name server
  interoperability 31-8

P

PortChannels
  interoperability 31-7
protocol analysis 31-2

R

remote capture 31-4
remote capture daemon 31-3
Remote Capture Protocol
  See RPCAP
RPCAP
  Ethereal communication 31-4

T

TE ports
  interoperability 31-7
TOV
  interoperability 31-6
  ranges 31-1
troubleshooting
  collecting output 31-8
trunking
  interoperability 31-7

V

VSANs
  interop mode 31-7

W

world wide names
  See See WWNs
WWNs
  configuring 31-5