



Send comments to nx5000-docfeedback@cisco.com



Fabric Manager Software Configuration Guide for the Cisco Nexus 5000 Series Switch

Fabric Manager Software Release 3.4(1)

June 2008

Americas 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: 408 527-0883

Text Part Number: OL-16598-01

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

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

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

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

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

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

Fabric Manager Software Configuration Guide for the Cisco Nexus 5000 Series Switch
© 2008 Cisco Systems, Inc. All rights reserved

Send comments to nx5000-docfeedback@cisco.com



CONTENTS

Preface i

CHAPTER 1

Product Overview	1-1
New Technologies in the Cisco Nexus 5000 Series	1-1
Fibre Channel over Ethernet	1-1
I/O Consolidation	1-2
Virtual Interfaces	1-3
Cisco Nexus 5000 Series Switch Hardware	1-3
Chassis	1-3
Expansion Modules	1-3
Ethernet Interfaces	1-4
Fibre Channel Interfaces	1-4
Management Interfaces	1-4
Cisco Nexus 5000 Series Switch Software	1-4
Ethernet Switching	1-4
FCoE and Fibre Channel Switching	1-5
Licensing	1-5
QoS	1-5
Serviceability	1-5
Switched Port Analyzer	1-5
Ethanalyzer	1-6
Call Home	1-6
Online Diagnostics	1-6
Switch Management	1-6
Simple Network Management Protocol	1-6
Configuration Verification and Rollback	1-7
Role-Based Access Control	1-7
Configuration Methods	1-7
Network Security Features	1-7
Virtual Device Contexts	1-8
Typical Deployment Topologies	1-8
Ethernet TOR Switch Topology	1-8
IOC Topology	1-10
Supported Standards	1-11

Send comments to nx5000-docfeedback@cisco.com

CHAPTER 2

Installing Cisco Fabric Manager 2-1

- Information About Cisco Fabric Manager 2-1
 - Fabric Manager Server 2-2
 - Fabric Manager Client 2-2
 - Fabric Manager Server Proxy Services 2-2
 - Device Manager 2-3
 - Performance Manager 2-3
 - Fabric Manager Web Server 2-3
- Understanding Switch Management 2-3
 - mgmt0 2-4
 - IPFC 2-5
- Installing the Management Software 2-5
 - Before You Install 2-5
 - Supported Software 2-6
 - Minimum Hardware Requirements 2-7
 - Installing the Database 2-7
 - Installing Oracle 2-7
 - Installing Fabric Manager 2-8
 - Installing Device Manager 2-16
- Upgrading the Management Software 2-18
- Integrating Cisco Fabric Manager with Other Management Tools 2-18
- Running Fabric Manager Behind a Firewall 2-19
- Uninstalling the Management Software 2-20

CHAPTER 3

Fabric Manager Server 3-1

- Information About Fabric Manager Server 3-1
- Installing and Configuring Fabric Manager Server 3-2
 - Installing Fabric Manager Server 3-2
 - Unlicensed Versus Licensed Fabric Manager Server 3-3
 - Installing Fabric Manager Web Server 3-3
 - Verifying Performance Manager Collections 3-3
- Managing a Fabric Manager Server Fabric 3-3
 - Selecting a Fabric to Manage Continuously 3-3
- Fabric Manager Server Properties File 3-4
- Modifying Fabric Manager Server 3-5
 - Adding or Removing Fabric Manager Server Users 3-5
 - Changing the Fabric Manager Server User Name and Password 3-6
 - Changing the Polling Period and Fabric Rediscovery Time 3-6

Send comments to nx5000-docfeedback@cisco.com

- Using Device Aliases or FC Aliases 3-7
- Saving Device Aliases to the Switch 3-7

 CHAPTER 4

- Authentication in Fabric Manager 4-1**
 - Information About Fabric Manager Authentication 4-1
 - Discovering a Fabric 4-3
 - Setting Up Discovery for a Fabric 4-3
 - Authenticating Performance Manager 4-4
 - Authenticating Fabric Manager Web Server 4-5

 CHAPTER 5

- Fabric Manager Client 5-1**
 - Information About Fabric Manager Client 5-1
 - Fabric Manager Advanced Mode 5-1
 - Launching Fabric Manager Client 5-2
 - Fabric Manager Client Quick Tour 5-6
 - Menu Bar 5-7
 - Toolbar 5-8
 - Logical Domains Pane 5-9
 - Filtering 5-10
 - Physical Attributes Pane 5-10
 - Information Pane 5-11
 - Detachable Tables 5-12
 - Fabric Pane 5-12
 - Context Menus 5-14
 - Saving the Map 5-14
 - Purging Down Elements 5-15
 - Multiple Fabric Display 5-15
 - Filtering by Groups 5-16
 - Status Bar 5-18
 - Setting Fabric Manager Preferences 5-18
 - Network Fabric Discovery 5-19
 - Modifying the Device Grouping 5-20
 - Using Alias Names as Enclosures 5-20
 - Controlling Administrator Access with Users and Roles 5-21
 - Using Fabric Manager Wizards 5-21
 - Fabric Manager Troubleshooting Tools 5-22

Send comments to nx5000-docfeedback@cisco.com

CHAPTER 6

Device Manager 6-1

- Information About Device Manager 6-1
- Launching Device Manager 6-2
- Using Device Manager 6-2
 - Menu Bar 6-2
 - Toolbar Icons 6-3
 - Dialog Boxes 6-4
 - Tabs 6-4
 - Legend 6-4
 - Supervisor and Switching Modules 6-5
 - Context Menus 6-6
- Using the Quick Configuration Tool 6-6
- Setting Device Manager Preferences 6-7

CHAPTER 7

Using Cisco Fabric Services 7-1

- Information About CFS 7-1
- CFS Distribution 7-2
 - CFS Distribution Modes 7-2
 - Uncoordinated Distribution 7-3
 - Coordinated Distribution 7-3
 - Unrestricted Uncoordinated Distributions 7-3
 - Enabling/Disabling CFS Distribution on a Switch 7-3
 - CFS Distribution over IP 7-4
 - CFS Distribution over Fibre Channel 7-5
 - CFS Distribution Scopes 7-5
 - CFS Merge Support 7-6
- CFS Support for Applications 7-6
 - CFS Application Requirements 7-6
 - Enabling CFS for an Application 7-7
 - Locking the Network 7-8
 - Committing Changes 7-8
 - Discarding Changes 7-9
 - Saving the Configuration 7-10
 - Clearing a Locked Session 7-10
- CFS Regions 7-10
 - About CFS Regions 7-11
 - Example Scenario 7-11
 - Managing CFS Regions Using Fabric Manager 7-11

Send comments to nx5000-docfeedback@cisco.com

Creating CFS Regions	7-12
Assigning Features to CFS Regions	7-12
Moving a Feature to a Different Region	7-13
Removing a Feature from a Region	7-14
Deleting CFS Regions	7-14
Displaying CFS Distribution Information	7-15
CFS Example Using Fabric Manager	7-15
CFS Example Using Device Manager	7-18
Default Settings	7-19

 CHAPTER 8

Configuring Ethernet Interfaces	8-1
Information About Ethernet Interfaces	8-1
Configuring Ethernet Interfaces	8-1
Displaying Interface Information	8-1
Default Physical Ethernet Settings	8-2

 CHAPTER 9

Configuring Virtual Interfaces	9-1
Information About Virtual Interfaces	9-1
Configuring Virtual Interfaces	9-1
Creating a Virtual Interface Group	9-2
Using the Virtual Interface Group Wizard	9-3
Binding a VIG to a Physical Ethernet Interface	9-4
Deleting a Virtual Interface Group	9-4
Using the Virtual Interface Wizard	9-5
Creating a Virtual Ethernet Interface	9-6
Deleting a Virtual Ethernet Interface	9-7
Creating a Virtual Fibre Channel Interface	9-7
Deleting a Virtual Fibre Channel Interface	9-8

 CHAPTER 10

Configuring Fibre Channel Interfaces	10-1
Information About Fibre Channel Interfaces	10-1
Licensing Requirements	10-1
Physical Fibre Channel Interfaces	10-1
Virtual Fibre Channel Interfaces	10-2
Interface Modes	10-2
E Port	10-3
F Port	10-4
NP Port	10-4

Send comments to nx5000-docfeedback@cisco.com

- TE Port 10-4
- SD Port 10-4
- Auto Mode 10-4
- Interface States 10-5
 - Administrative States 10-5
 - Operational States 10-5
 - Reason Codes 10-5
- Buffer-to-Buffer Credits 10-7
- Configuring Fibre Channel Interfaces 10-8
 - Configuring a Fibre Channel Interface 10-8
 - Setting the Interface Administrative State 10-9
 - Configuring Interface Modes 10-9
 - Configuring the Interface Description 10-9
 - Configuring Administrative Speeds 10-10
 - Autosensing 10-10
 - Configuring SD Port Frame Encapsulation 10-10
 - Configuring Receive Data Field Size 10-11
 - Understanding Bit Error Thresholds 10-11
 - Configuring Buffer-to-Buffer Credits 10-12
- Verifying Fibre Channel Interfaces 10-12
 - Verifying SFP Transmitter Types 10-13
 - Obtaining Interface Statistics 10-13
- Default Settings 10-14

CHAPTER 11

- Configuring Domain Parameters 11-1**
 - Information About Fibre Channel Domains 11-1
 - About Domain Restart 11-3
 - Restarting a Domain 11-3
 - About Switch Priority 11-4
 - Configuring Switch Priority 11-4
 - About fcdomain Initiation 11-5
 - Enabling or Disabling fcdomains 11-5
 - Setting Fabric Names 11-6
 - About Incoming RCFs 11-6
 - Rejecting Incoming RCFs 11-6
 - About Autoreconfiguring Merged Fabrics 11-7
 - Enabling Autoreconfiguration 11-7
 - Domain IDs 11-8
 - About Domain IDs 11-8

Send comments to nx5000-docfeedback@cisco.com

Specifying Static or Preferred Domain IDs	11-10
About Allowed Domain ID Lists	11-10
Configuring Allowed Domain ID Lists	11-11
About CFS Distribution of Allowed Domain ID Lists	11-12
Enabling Distribution	11-12
Locking the Fabric	11-12
Committing Changes	11-12
Discarding Changes	11-13
Clearing a Fabric Lock	11-13
Displaying Pending Changes	11-14
Displaying Session Status	11-14
About Contiguous Domain ID Assignments	11-14
Enabling Contiguous Domain ID Assignments	11-15
FC IDs	11-15
About Persistent FC IDs	11-16
Enabling the Persistent FC ID Feature	11-16
Persistent FC ID Configuration Guidelines	11-16
Configuring Persistent FC IDs	11-17
About Unique Area FC IDs for HBAs	11-17
Configuring Unique Area FC IDs for an HBA	11-18
About Persistent FC ID Selective Purging	11-19
Purging Persistent FC IDs	11-19
Displaying fcdomain Statistics	11-20
Default Settings	11-21

CHAPTER 12
Configuring N-Port Virtualization 12-1

Information About NPV	12-1
NP Ports	12-2
NP Links	12-2
FLOGI Operation	12-2
Guidelines and Limitations	12-4
Configuring NPV	12-4
Configuring NPV with Device Manager	12-5

CHAPTER 13
Configuring VSAN Trunking 13-1

Information About VSAN Trunking	13-1
VSAN Trunking Mismatches	13-2
VSAN Trunking Protocol	13-2
Configuring VSAN Trunking	13-3

Send comments to nx5000-docfeedback@cisco.com

- Guidelines and Restrictions 13-3
- About Trunk Mode 13-3
- Configuring Trunk Mode 13-4
- About Trunk-Allowed VSAN Lists 13-5
- Configuring an Allowed-Active List of VSANs 13-6
- Default Settings 13-7

CHAPTER 14

- Configuring SAN Port Channels 14-1**
 - Information About SAN Port Channels 14-1
 - Understanding Port Channels and VSAN Trunking 14-2
 - Understanding Load Balancing 14-3
 - Configuring SAN Port Channels 14-5
 - SAN Port Channel Configuration Guidelines 14-6
 - Configuring SAN Port Channels 14-7
 - About SAN Port Channel Modes 14-10
 - About SAN Port Channel Deletion 14-11
 - Deleting SAN Port Channels 14-11
 - Interfaces in a SAN Port Channel 14-12
 - About Interface Addition to a SAN Port Channel 14-12
 - Compatibility Check 14-12
 - Suspended and Isolated States 14-13
 - Adding an Interface to a SAN Port Channel 14-13
 - Forcing an Interface Addition 14-14
 - About Interface Deletion from a SAN Port Channel 14-14
 - Deleting an Interface from a SAN Port Channel 14-14
 - Port Channel Protocol 14-15
 - About Channel Group Creation 14-15
 - Autocreation Guidelines 14-17
 - Enabling and Configuring Autocreation 14-17
 - About Manually Configured Channel Groups 14-18
 - Converting to Manually Configured Channel Groups 14-18
 - Verifying SAN Port Channel Configuration 14-19
 - Default Settings 14-19

CHAPTER 15

- Configuring and Managing VSANs 15-1**
 - Information About VSANs 15-1
 - VSAN Topologies 15-1
 - VSAN Advantages 15-4
 - VSANs Versus Zones 15-4

Send comments to nx5000-docfeedback@cisco.com

Configuring VSANs	15-5
About VSAN Creation	15-6
Creating VSANs Statically	15-6
About Port VSAN Membership	15-8
Assigning Static Port VSAN Membership	15-8
About the Default VSAN	15-8
About the Isolated VSAN	15-8
Displaying Isolated VSAN Membership	15-9
Operational State of a VSAN	15-9
About Static VSAN Deletion	15-9
Deleting Static VSANs	15-10
About Load Balancing	15-11
Configuring Load Balancing	15-11
About Interop Mode	15-12
Default Settings	15-12

CHAPTER 16

Configuring and Managing Zones	16-1
Information About Zoning	16-1
Zoning Features	16-2
Zoning Example	16-3
Zone Implementation	16-4
Active and Full Zone Set Configuration Guidelines	16-5
Configuring Zones	16-7
About the Zone Configuration Tool	16-7
Configuring Zones Using the Zone Configuration Tool	16-8
Adding Zone Members	16-10
Configuring the Default Zone Policy	16-12
Zone Sets	16-12
About Zone Set Creation	16-13
Activating a Zone Set	16-13
Displaying Zone Membership Information	16-15
About the Default Zone	16-16
Configuring the Default Zone	16-16
About FC Alias Creation	16-17
Creating FC Aliases	16-17
Adding Members to Aliases	16-18
Converting Zone Members to pWWN-Based Members	16-20
Zone Enforcement	16-21
Zone Set Distribution	16-21

Send comments to nx5000-docfeedback@cisco.com

- Enabling Full Zone Set Distribution 16-21
- Enabling a One-Time Distribution 16-22
- About Recovering from Link Isolation 16-23
- Importing and Exporting Zone Sets 16-23
- Zone Set Duplication 16-24
 - Copying Zone Sets 16-25
 - About Backing Up and Restoring Zones 16-25
 - Backing Up and Restoring Zones 16-25
 - Renaming Zones, Zone Sets, and Aliases 16-26
 - Cloning Zones, Zone Sets, FC Aliases, and Zone Attribute Groups 16-27
 - Migrating a Non-MDS Database 16-28
 - Clearing the Zone Server Database 16-28
- Verifying Zone Information 16-28
- Enhanced Zoning 16-29
 - About Enhanced Zoning 16-29
 - Changing from Basic Zoning to Enhanced Zoning 16-30
 - Changing from Enhanced Zoning to Basic Zoning 16-30
 - Enabling Enhanced Zoning 16-31
 - Merging the Database 16-31
 - Analyzing a Zone Merge 16-32
 - Configuring Zone Merge Control Policies 16-32
- Compacting the Zone Database 16-33
- Default Settings 16-33

CHAPTER 17

- Distributing Device Alias Services 17-1**
 - Information About Device Aliases 17-1
 - Device Alias Features 17-1
 - Device Alias Requirements 17-2
 - Zone Aliases Versus Device Aliases 17-2
 - Device Alias Databases 17-3
 - Device Alias Modes 17-3
 - Changing Device Alias Mode Guidelines 17-3
 - Configuring Device Alias Modes 17-4
 - About Device Alias Distribution 17-5
 - Distributing the Device Alias Database 17-5
 - About Creating a Device Alias 17-5
 - Creating a Device Alias 17-6
 - Committing Changes 17-6
 - Discarding Changes 17-7

Send comments to nx5000-docfeedback@cisco.com

Legacy Zone Alias Conversion	17-7
Using Device Aliases or FC Aliases	17-7
Database Merge Guidelines	17-8
Default Settings	17-9

CHAPTER 18
Configuring Fibre Channel Routing Services and Protocols 18-1

Information About FSPF	18-1
FSPF Examples	18-2
Fault Tolerant Fabric Example	18-2
Redundant Link Example	18-2
FSPF Global Configuration	18-3
About SPF Computational Hold Times	18-3
About Link State Records	18-3
Configuring FSPF on a VSAN	18-4
Resetting FSPF to the Default Configuration	18-5
Enabling or Disabling FSPF	18-5
FSPF Interface Configuration	18-5
About FSPF Link Cost	18-6
Configuring FSPF Link Cost	18-6
About Hello Time Intervals	18-6
Configuring Hello Time Intervals	18-7
About Dead Time Intervals	18-7
Configuring Dead Time Intervals	18-7
About Retransmitting Intervals	18-7
Configuring Retransmitting Intervals	18-8
About Disabling FSPF for Specific Interfaces	18-8
Disabling FSPF for Specific Interfaces	18-8
Displaying the FSPF Database	18-9
Viewing FSPF Statistics	18-10
FSPF Routes	18-11
About Fibre Channel Routes	18-11
Configuring Fibre Channel Routes	18-11
In-Order Delivery	18-12
About Reordering Network Frames	18-13
About Reordering SAN Port Channel Frames	18-13
About Enabling In-Order Delivery	18-14
Enabling In-Order Delivery Globally	18-15
Enabling In-Order Delivery for a VSAN	18-15
Configuring the Drop Latency Time	18-15

Send comments to nx5000-docfeedback@cisco.com

Default Settings 18-16

CHAPTER 19

Managing FLOGI, Name Server, FDMI, and RSCN Databases 19-1

Information About Fabric Login 19-1

Name Server Proxy 19-2

About Registering Name Server Proxies 19-2

Registering Name Server Proxies 19-2

About Rejecting Duplicate pWWNs 19-3

Rejecting Duplicate pWWNs 19-3

About Name Server Database Entries 19-3

Viewing Name Server Database Entries 19-3

FDMI 19-4

Displaying FDMI 19-4

RSCN 19-5

About RSCN Information 19-5

Displaying RSCN Information 19-5

About the multi-pid Option 19-6

Configuring the multi-pid Option 19-6

Clearing RSCN Statistics 19-7

RSCN Timer Configuration Distribution Using CFS 19-7

Configuring the RSCN Timer with CFS 19-8

Default Settings 19-8

CHAPTER 20

Configuring SPAN 20-1

Information About SPAN Sources 20-1

Characteristics of Source Ports 20-1

Information About SPAN Destinations 20-2

Characteristics of Destination Ports 20-2

Configuring SPAN 20-3

Configuring SPAN Using Device Manager 20-3

Creating SPAN Sessions Using Device Manager 20-3

Editing SPAN Sources Using Device Manager 20-4

Deleting SPAN Sessions Using Device Manager 20-5

Default SPAN Settings 20-5

CHAPTER 21

Discovering SCSI Targets 21-1

Information About SCSI LUN Discovery 21-1

About Starting SCSI LUN Discovery 21-1

Send comments to nx5000-docfeedback@cisco.com

Starting SCSI LUN Discovery	21-2
About Initiating Customized Discovery	21-2
Initiating Customized Discovery	21-2
Displaying SCSI LUN Information	21-3

 CHAPTER 22

Advanced Features and Concepts 22-1

Fibre Channel Timeout Values	22-1
Timer Configuration Across All VSANs	22-2
Timer Configuration Per-VSAN	22-3
About fctimer Distribution	22-4
Enabling or Disabling fctimer Distribution	22-4
Database Merge Guidelines	22-4
World Wide Names	22-5
Verifying WWN Information	22-5
Link Initialization WWN Usage	22-5
Configuring a Secondary MAC Address	22-6
FC ID Allocation for HBAs	22-6
Default Company ID List	22-7
Verifying the Company ID Configuration	22-7
Switch Interoperability	22-7
About Interop Mode	22-8
Configuring Interop Mode 1	22-9
Verifying Interoperating Status	22-11
Default Settings	22-12

 CHAPTER 23

Configuring FC-SP and DHCHAP 23-1

Information About Fabric Authentication	23-1
DHCHAP	23-2
DHCHAP Compatibility with Fibre Channel Features	23-3
About Enabling DHCHAP	23-4
Enabling DHCHAP	23-4
About DHCHAP Authentication Modes	23-4
Configuring the DHCHAP Mode	23-5
About the DHCHAP Hash Algorithm	23-6
Configuring the DHCHAP Hash Algorithm	23-6
About the DHCHAP Group Settings	23-6
Configuring the DHCHAP Group Settings	23-6
About the DHCHAP Password	23-7
Configuring DHCHAP Passwords for the Local Switch	23-7

Send comments to nx5000-docfeedback@cisco.com

- About Password Configuration for Remote Devices 23-8
- Configuring DHCHAP Passwords for Remote Devices 23-8
- About the DHCHAP Timeout Value 23-8
- Configuring the DHCHAP Timeout Value 23-9
- Configuring DHCHAP AAA Authentication 23-9
- Enabling FC-SP on ISLs 23-9
- Default Settings 23-10

CHAPTER 24

Configuring Port Security 24-1

- Information About Port Security 24-1
 - Port Security Enforcement 24-2
 - About Auto-Learning 24-2
 - Port Security Activation 24-3
- Configuring Port Security 24-3
 - Configuring Port Security with Auto-Learning and CFS Distribution 24-3
 - Configuring Port Security with Auto-Learning without CFS 24-4
 - Configuring Port Security with Manual Database Configuration 24-5
- Enabling Port Security 24-5
- Port Security Activation 24-6
 - Activating Port Security 24-7
 - Database Activation Rejection 24-7
 - Forcing Port Security Activation 24-8
 - Database Reactivation 24-8
 - Copying an Active Database to the Config Database 24-9
 - Displaying Activated Port Security Settings 24-9
 - Displaying Port Security Statistics 24-9
 - Displaying Port Security Violations 24-10
- Auto-Learning 24-10
 - About Enabling Auto-Learning 24-10
 - Enabling Auto-Learning 24-11
 - Disabling Auto-Learning 24-11
 - Auto-Learning Device Authorization 24-12
 - Authorization Scenario 24-12
- Port Security Manual Configuration 24-13
 - WWN Identification Guidelines 24-14
 - Adding Authorized Port Pairs 24-14
 - Deleting Port Security Setting 24-15
- Port Security Configuration Distribution 24-15
 - Enabling Distribution 24-16

Send comments to nx5000-docfeedback@cisco.com

Locking the Fabric	24-16
Committing the Changes	24-17
Activation and Auto-Learning Configuration Distribution	24-17
Database Merge Guidelines	24-18
Database Interaction	24-18
Database Scenarios	24-19
Copying the Port Security Database	24-20
Deleting the Port Security Database	24-20
Clearing the Port Security Database	24-21
Default Settings	24-21

 CHAPTER 25

Configuring Fabric Binding	25-1
Information About Fabric Binding	25-1
Licensing Requirements	25-1
Port Security Versus Fabric Binding	25-2
Fabric Binding Enforcement	25-2
Configuring Fabric Binding	25-3
Configuring Fabric Binding	25-3
Enabling Fabric Binding	25-4
About Switch WWN Lists	25-4
Configuring Switch WWN List	25-4
About Fabric Binding Activation and Deactivation	25-5
Activating Fabric Binding	25-5
Forcing Fabric Binding Activation	25-6
Copying Fabric Binding Configurations	25-6
Creating a Fabric Binding Configuration	25-7
Deleting a Fabric Binding Configuration	25-7
Copying Fabric Binding to the Configuration File	25-8
Viewing EFMD Statistics	25-8
Viewing Fabric Binding Violations	25-8
Viewing Fabric Binding Active Database	25-8
Saving Fabric Binding Configurations	25-9
Clearing the Fabric Binding Statistics	25-9
Deleting the Fabric Binding Database	25-10
Default Settings	25-10

 CHAPTER 26

Configuring Fabric Configuration Servers	26-1
Information About FCS	26-1
FCS Characteristics	26-2

Send comments to nx5000-docfeedback@cisco.com

- Displaying FCS Discovery 26-3
- Displaying FCS Elements 26-3
- Creating an FCS Platform 26-4
- Displaying FCS Fabric Ports 26-5
- Default Settings 26-6

CHAPTER 27

- Configuring Port Tracking 27-1**
 - Information About Port Tracking 27-1
 - Configuring Port Tracking 27-2
 - Enabling Port Tracking 27-3
 - About Configuring Linked Ports 27-3
 - Operationally Binding a Tracked Port 27-4
 - About Tracking Multiple Ports 27-5
 - Tracking Multiple Ports 27-6
 - About Monitoring Ports in a VSAN 27-6
 - Monitoring Ports in a VSAN 27-6
 - About Forceful Shutdown 27-6
 - Forcefully Shutting Down a Tracked Port 27-6
 - Default Port Tracking Settings 27-7

CHAPTER 28

- Network Monitoring 28-1**
 - Information About SAN Discovery and Topology Mapping 28-1
 - Device Discovery 28-1
 - Topology Mapping 28-2
 - Using the Topology Map 28-2
 - Saving a Customized Topology Map Layout 28-2
 - Using Enclosures with Fabric Manager Topology Maps 28-3
 - Mapping Multiple Fabrics 28-3
 - Inventory Management 28-3
 - Using the Inventory Tab from Fabric Manager Web Server 28-4
 - Viewing Logs from Device Manager 28-4
 - Health and Event Monitoring 28-4
 - Fabric Manager Events Tab 28-4
 - Event Information in Fabric Manager Web Server Reports 28-5
 - Events in Device Manager 28-5

CHAPTER 29

- Performance Manager 29-1**
 - Information About Performance Manager 29-1

Send comments to nx5000-docfeedback@cisco.com

Data Interpolation	29-2
Data Collection	29-2
Using Performance Thresholds	29-2
Flow Setup Wizards	29-3
Flow Statistics Configuration	29-4
About Flow Statistics	29-4
Counting Flow Statistics	29-4

 CHAPTER 30

Nexus 5000 Management Software FAQ 30-1

Nexus 5000 Series Issues	30-1
What is Display FCoE Mode?	30-1
Switching to Display FCoE Mode	30-1
General Fabric Manager Issues	30-2

 CHAPTER 31

Troubleshooting Your Fabric 31-1

Troubleshooting Tools and Techniques	31-1
Cisco Traffic Analyzer	31-2
Cisco Protocol Analyzer	31-3
Analyzing Switch Device Health	31-3
Analyzing Switch Fabric Configuration	31-4
Analyzing End-to-End Connectivity	31-5
Using the Ping Tool (fcping)	31-7
Using Trace Route (fctrace) and Other Troubleshooting Tools	31-7
Analyzing the Results of Merging Zones	31-8
Using the Show Tech Support Command	31-9
Running CLI Commands	31-11
Adjusting for Daylight Savings Time	31-12
Locating Other Switches	31-12
Fibre Channel Timeout Values	31-14
Timer Configuration Per-VSAN	31-15
Configuring a Fabric Analyzer	31-16
About the Cisco Fabric Analyzer	31-17
Local Text-Based Capture	31-17
Remote Capture Daemon	31-18
GUI-Based Client	31-18
Configuring the Cisco Fabric Analyzer	31-18
Sending Captures to Remote IP Addresses	31-19
Displaying Captured Frames	31-19

Send comments to nx5000-docfeedback@cisco.com

- Defining Display Filters 31-20
 - Capture Filters 31-20
 - Permitted Capture Filters 31-21
- Configuring World Wide Names 31-22
 - Link Initialization WWN Usage 31-22
- Configuring a Secondary MAC Address 31-22
 - Displaying WWN Information 31-23
- FC ID Allocation for HBAs 31-23
- Default Settings 31-23

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)



Preface

This preface describes the audience, organization, and conventions of the *Cisco Nexus 5000 Series Fabric Manager Configuration Guide*. It also provides information on how to obtain related documentation.

Audience

This guide is for experienced network administrators who are responsible for configuring and maintaining Cisco Nexus 5000 Series switches.

Organization

This guide is organized as follows:

Chapter	Title	Description
Chapter 1	Product Overview	Provides an overview of Cisco Nexus 5000 Series switches.
Chapter 2	Installing Cisco Fabric Manager	Provides a brief overview of Fabric Manager components and capabilities, and information on installation and launching the applications.
Chapter 3	Fabric Manager Server	Provides in-depth descriptions of GUI and capabilities for the Fabric Manager Server.
Chapter 4	Authentication in Fabric Manager	Describes the authentication schemes between Fabric Manager components and fabric switches.
Chapter 5	Fabric Manager Client	Provides in-depth descriptions of GUI and capabilities for the Fabric Manager Client.
Chapter 6	Device Manager	Provides in-depth descriptions of GUI and capabilities for the Device Manager.
Chapter 7	Using Cisco Fabric Services	Provides descriptions of GUI and capabilities for Cisco Fabric Services.
Chapter 8	Configuring Ethernet Interfaces	Provides descriptions of how to configure Ethernet interfaces using Fabric Manager.

Send comments to nx5000-docfeedback@cisco.com

Chapter	Title	Description
Chapter 9	Configuring Virtual Interfaces	Provides descriptions of how to configure virtual interfaces using Fabric Manager.
Chapter 10	Configuring Fibre Channel Interfaces	Provides descriptions of how to configure Fibre Channel interfaces using Fabric Manager.
Chapter 11	Configuring Domain Parameters	Provides descriptions of how to configure Fibre Channel domains using Fabric Manager.
Chapter 12	Configuring N-Port Virtualization	Explains how to configure NPV devices to reduce excessive Fibre Channel domain IDs in SANs.
Chapter 13	Configuring VSAN Trunking	Provides descriptions of how to configure VSAN trunks using Fabric Manager.
Chapter 14	Configuring SAN Port Channels	Explains SAN port channels and load balancing concepts and provides details on configuring SAN port channels, adding ports to SAN port channels, and deleting ports from SAN port channels.
Chapter 15	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 16	Configuring and Managing Zones	Defines various zoning concepts and provides details on configuring a zone set and zone management features.
Chapter 17	Distributing Device Alias Services	Describes the use of the Distributed Device Alias Services (device alias) to distribute device alias names on a fabric-wide basis.
Chapter 18	Configuring Fibre Channel Routing Services and Protocols	Provides details and configuration information on Fibre Channel routing services and protocols.
Chapter 19	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 20	Configuring SPAN	Describes the Switched Port Analyzer (SPAN), SPAN sources, filters, SPAN sessions, SD port characteristics, and configuration details.
Chapter 21	Discovering SCSI Targets	Describes how the SCSI LUN discovery feature is started and displayed.
Chapter 14	Configuring SAN Port Channels	Explains SAN port channels and load balancing concepts and provides details on configuring SAN port channels, adding ports to SAN port channels, and deleting ports from SAN port channels.

Send comments to nx5000-docfeedback@cisco.com

Chapter	Title	Description
Chapter 22	Advanced Features and Concepts	Describes the advanced configuration features, including timeout values, the fctrace tool, and interoperating with non-Cisco switches.
Chapter 23	Configuring FC-SP and DHCHAP	Describes the DHCHAP protocol, which is an FC-SP protocol to provide authentication between Cisco Nexus 5000 Series switches and other devices.
Chapter 24	Configuring Port Security	Provides details on port security features that can prevent unauthorized access to a switch port in a Cisco Nexus 5000 Series switch.
Chapter 25	Configuring Fabric Binding	Describes how to configure fabric binding capabilities using Fabric Manager.
Chapter 26	Configuring Fabric Configuration Servers	Describes how to configure fabric configuration servers using Fabric Manager.
Chapter 27	Configuring Port Tracking	Provides information about a port tracking feature that provides a faster recovery from link failures.
Chapter 28	Network Monitoring	Describes how to use Fabric Manager monitoring features.
Chapter 29	Performance Manager	Provides details on using Performance Manager.
Chapter 30	Nexus 5000 Management Software FAQ	Provides answers to frequently asked questions.
Chapter 31	Troubleshooting Your Fabric	Provides details on troubleshooting Fabric Manager.

Document Conventions

Command descriptions use these conventions:

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

Screen examples use these conventions:

<code>screen font</code>	Terminal sessions and information the switch displays are in screen font.
boldface screen font	Information you must enter is in boldface screen font.
<i>italic screen font</i>	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.

Send comments to nx5000-docfeedback@cisco.com

[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:



Note

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



Caution

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

Related Documentation

Documentation for Cisco Cisco Nexus 5000 Series switches is available at the following URL:

http://www.cisco.com/en/US/products/ps9670/tsd_products_support_series_home.html

The following are related Cisco Nexus 5000 Series documents:

Cisco Nexus 5000 Series Release Notes

Cisco Nexus 5000 Series CLI Software Configuration Guide, Release 4.0

Cisco Nexus 5000 Series Fabric Manager Software Configuration Guide, Release 4.0

Cisco Nexus 5000 Series System Messages Reference

Cisco Nexus 5000 Series Command Reference, Release 4.0

Cisco Nexus 5000 Series Hardware Installation Guide, Release 4.0

Cisco Nexus 5000 Series MIBs Reference, Release 4.0

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

<http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.



CHAPTER 1

Product Overview

The Cisco Nexus 5000 Series is a family of top-of-rack switches for the data center. The Nexus 5020 switch is a 10-Gigabit Ethernet and Fibre Channel over Ethernet (FCoE) switch with 1.04 Tbps switching throughput. The Nexus 5020 provides low-latency wire-speed switching for up to 52 10-Gigabit Ethernet ports.

The Nexus 5020 switch supports FCoE to provide data center I/O consolidation (IOC). Optional Fibre Channel-capable expansion modules provide four or eight native Fibre Channel SAN interfaces.

This chapter describes the Cisco Nexus 5000 Series switches and includes the following sections:

- [New Technologies in the Cisco Nexus 5000 Series, page 1-1](#)
- [Cisco Nexus 5000 Series Switch Hardware, page 1-3](#)
- [Cisco Nexus 5000 Series Switch Software, page 1-4](#)
- [Typical Deployment Topologies, page 1-8](#)
- [Supported Standards, page 1-11](#)

New Technologies in the Cisco Nexus 5000 Series

Cisco Nexus 5000 Series switches introduce several new technologies, which are described in the following sections:

- [Fibre Channel over Ethernet, page 1-1](#)
- [I/O Consolidation, page 1-2](#)
- [Virtual Interfaces, page 1-3](#)

Fibre Channel over Ethernet

Fibre Channel over Ethernet (FCoE) provides a method of encapsulating Fibre Channel traffic over a physical Ethernet link. FCoE frames use a unique Ethertype so that FCoE traffic and standard Ethernet traffic can be carried on the same link.

Fibre Channel traffic requires a lossless transport layer. Native Fibre Channel implements lossless service using a buffer-to-buffer credit system. For FCoE traffic, the Ethernet link must provide lossless service.

Ethernet links on Cisco Nexus 5000 Series switches provide two mechanisms to ensure lossless transport for FCoE traffic: link-level flow control and priority flow control.

Send comments to nx5000-docfeedback@cisco.com

IEEE 802.3x link-level flow control allows a congested receiver to signal the far end to pause the data transmission for a short period of time. The pause functionality is applied to all the traffic on the link.

The priority flow control (PFC) feature applies pause functionality to specific classes of traffic on the Ethernet link. For example, PFC can provide lossless service for the FCoE traffic, and best-effort service for the standard Ethernet traffic. PFC can provide different levels of service to specific classes of Ethernet traffic (using IEEE 802.1p traffic classes).

I/O Consolidation

I/O consolidation (IOC) allows a single network technology to carry IP, SAN and IPC traffic.


FCoE enables an evolutionary approach to IOC. The upper Fibre Channel layers are unchanged, so the Fibre Channel operational model is maintained. FCoE network management and configuration is similar to a native Fibre Channel network.

Cisco Nexus 5000 Series switches use FCoE to carry Fibre Channel and Ethernet traffic on the same physical Ethernet connection between the switch and the server. At the server, the connection terminates to a converged network adapter (CNA). The adapter presents two interfaces to the server's operating system (OS): one Ethernet NIC interface and one Fibre Channel HBA interface. The server OS is not aware of the FCoE encapsulation (See [Figure 1-1](#))

At the switch, the incoming Ethernet port separates the Ethernet and Fibre Channel traffic (using Ethertype to differentiate the frames). Ethernet frames and Fibre Channel frames are switched to their respective network-side interfaces.

Cisco Nexus 5000 Series switches provide quality of service (QoS) capabilities to ensure lossless service across the switch for Fibre Channel traffic. Best-effort service can be applied to all of the Ethernet traffic or specific classes of Ethernet traffic can be configured with different QoS levels.

Figure 1-1 I/O Consolidation



[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Virtual Interfaces

When FCoE is enabled, a physical Ethernet cable carries traffic for a logical Ethernet connection and a logical Fibre Channel connection.

The Cisco Nexus 5000 Series switch uses virtual interfaces to represent the logical connections that are carried on the same physical Ethernet. The Cisco Nexus 5000 Series switch supports virtual Ethernet and virtual Fibre Channel interfaces.

For configuration purposes, virtual Ethernet and virtual Fibre Channel interfaces are implemented as Layer 2 subinterfaces of the physical Ethernet interface.

Link-level features (such as link debounce timer and CDP) are configured on the physical Ethernet interface. Logical Layer 2 Ethernet features (such as VLAN membership and ACLs) are configured on the virtual Ethernet interfaces. Logical Fibre Channel features (such as VSAN membership) are configured on the virtual Fibre Channel interfaces.

Cisco Nexus 5000 Series Switch Hardware

The Cisco Nexus 5000 Series includes the Nexus 5020 switch. The Nexus 5020 switch hardware is described in the following topics:

- [Chassis, page 1-3](#)
- [Expansion Modules, page 1-3](#)
- [Ethernet Interfaces, page 1-4](#)
- [Fibre Channel Interfaces, page 1-4](#)
- [Management Interfaces, page 1-4](#)

Chassis

The Nexus 5020 switch is a 2 RU chassis designed for rack mounting. The chassis supports redundant fans and power supplies.

The Nexus 5020 switching fabric is low latency, nonblocking and supports Ethernet frame sizes from 64 to 9216 bytes.

Expansion Modules

The Nexus 5020 switch has two slots for optional expansion modules. The following expansion modules are available:

- N5K-M1404 provides four 10-Gigabit Ethernet ports, and four 1/2/4 Gb Fibre Channel ports.
- N5K-M1600 provides six 10-Gigabit Ethernet ports.

The expansion modules are field-replaceable units (FRUs) that support online insertion and removal (OIR).

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Ethernet Interfaces

The Nexus 5020 switch has 40 fixed 10-Gigabit Ethernet ports equipped with SFP+ interface adapters. Up to 12 additional 10-Gigabit Ethernet ports are available on the expansion modules.

All of the 10-Gigabit Ethernet ports support FCoE. Each port can be used as a downlink (connected to a server) or as an uplink (to the data center LAN).

Fibre Channel Interfaces

Fibre Channel ports are optional on the Nexus 5020 switch. Up to eight Fibre Channel ports are available when using expansion modules.

Each Fibre Channel port can be used as a downlink (connected to a server) or as an uplink (to the data center SAN fabric).

Management Interfaces

The Nexus 5020 switch has two dedicated management interfaces (one serial console port and one 10/100/1000 Ethernet interface).

Cisco Nexus 5000 Series Switch Software

The Cisco Nexus 5000 Series switch is a Layer 2 device, which runs the Cisco Nexus operating system (NX-OS). The Cisco Nexus 5000 Series switch software is described in the following topics:

- [Ethernet Switching, page 1-4](#)
- [FCoE and Fibre Channel Switching, page 1-5](#)
- [Licensing, page 1-5](#)
- [QoS, page 1-5](#)
- [Serviceability, page 1-5](#)
- [Switch Management, page 1-6](#)
- [Network Security Features, page 1-7](#)
- [Virtual Device Contexts, page 1-8](#)

Ethernet Switching

Cisco Nexus 5000 Series switches are designed to support high-density, high-performance Ethernet systems and provide the following Ethernet switching features:

- IEEE 802.1D-2004 Rapid and Multiple Spanning Tree Protocols (802.1w and 802.1s)
- IEEE 802.1Q VLANs and trunks
- IEEE 802.3ad link aggregation
- Private VLANs
- Traffic suppression (unicast, multicast, and broadcast)

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

FCoE and Fibre Channel Switching

Cisco Nexus 5000 Series switches support data center I/O consolidation (IOC) by providing FCoE interfaces (to the servers) and native Fibre Channel interfaces (to the SAN).

FCoE and Fibre Channel switching includes the following features:

- Cisco fabric services
- N-port virtualization
- VSANs and VSAN trunking
- Zoning
- Distributed device alias service
- SAN port channels

Licensing

Cisco Cisco Nexus 5000 Series switches are shipped with the licenses installed. The switch provides commands to manage the licenses and install additional licenses.

QoS

The Cisco Nexus 5000 Series switch provides quality of service (QoS) capabilities such as traffic prioritization and bandwidth allocation on egress interfaces.

The default QoS configuration on the switch provides lossless service for Fibre Channel and FCoE traffic. QoS can be configured to provide additional classes of service for Ethernet traffic.

Serviceability

The Cisco Nexus 5000 Series switch serviceability functions provide data for network planning and help to improve problem resolution time.

This section includes the following topics:

- [Switched Port Analyzer, page 1-5](#)
- [Ethanalyzer, page 1-6](#)
- [Call Home, page 1-6](#)
- [Online Diagnostics, page 1-6](#)
- [Embedded Event Manager, page 1-6](#)

Switched Port Analyzer

The switched port analyzer (SPAN) feature allows an administrator to analyze all traffic between ports by nonintrusively directing the SPAN session traffic to a SPAN destination port that has an external analyzer attached to it.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Ethalyzer

Ethalyzer is a Cisco NX-OS protocol analyzer tool based on the Wireshark (formerly Ethereal) open source code. Ethalyzer is a command-line version of Wireshark for capturing and decoding packets. You can use Ethalyzer to troubleshoot your network and analyze the control-plane traffic. For more information about Ethalyzer, see *Cisco NX-OS Troubleshooting Guide, Release 4.0*.

Call Home

The Call Home feature continuously monitors hardware and software components to provide e-mail-based notification of critical system events. A versatile range of message formats is available for optimal compatibility with pager services, standard e-mail, and XML-based automated parsing applications. The feature offers alert grouping capabilities and customizable destination profiles. This feature can be used, for example, to directly page a network support engineer, send an e-mail message to a network operations center (NOC), and employ Cisco AutoNotify services to directly generate a case with the Cisco Technical Assistance Center (TAC). This feature is a step toward autonomous system operation, which enables networking devices to inform IT when a problem occurs and helps to ensure that the problem is resolved quickly.

Online Diagnostics

Cisco generic online diagnostics (GOLD) is a suite of diagnostic facilities to verify that hardware and internal data paths are operating as designed. Boot-time diagnostics, continuous monitoring, and on-demand and scheduled tests are part of the Cisco GOLD feature set. GOLD allows rapid fault isolation and continuous system monitoring.

Switch Management

This section includes the following topics:

- [Simple Network Management Protocol, page 1-6](#)
- [Configuration Verification and Rollback, page 1-7](#)
- [Role-Based Access Control, page 1-7](#)
- [Configuration Methods, page 1-7](#)

Simple Network Management Protocol

Cisco NX-OS is compliant with Simple Network Management Protocol (SNMP) version 1, version 2, and version 3. A full set of Management Information Bases (MIBs) is supported.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Configuration Verification and Rollback

With the Cisco Nexus 5000 Series switch, you can verify the consistency of a configuration and the availability of necessary hardware resources before committing the configuration. A device can be preconfigured, and the verified configuration can be applied at a later time. Configurations also include checkpoints to allow the switch operator to revert to a known good configuration as needed.

Role-Based Access Control

With role-based access control (RBAC), you can limit access to switch operations by assigning roles to users. Administrators can customize access and restrict it to the users who require it.

Configuration Methods

You can configure Cisco Nexus 5000 Series switches using direct network configuration methods or web services hosted on a Fabric Manager server.

This section includes the following topics:

- [Configuring with CLI, XML Management Interface, or SNMP, page 1-7](#)
- [Configuring with Cisco MDS Fabric Manager, page 1-7](#)

Configuring with CLI, XML Management Interface, or SNMP

You can configure Cisco Nexus 5000 Series switches using the command line interface (CLI), the XML management interface over SSH, or SNMP as follows:

- CLI—You can configure switches using the CLI from an SSH session, a Telnet session, or the console port. SSH provides a secure connection to the device.
- XML Management Interface over SSH—You can configure switches using the XML management interface, which is a programming interface based on the NETCONF protocol that complements the CLI functionality. For more information, see the *Cisco NX-OS XML Management Interface User Guide, Release 4.0*.
- SNMP—SNMP allows you to configure switches using Management Information Bases (MIBs).

Configuring with Cisco MDS Fabric Manager

You can configure Cisco Nexus 5000 Series switches using the Fabric Manager client, which runs on a local PC and uses the Fabric Manager server.

Network Security Features

Cisco NX-OS Release 4.0 includes the following security features:

- Authentication, authorization, and accounting (AAA) and TACACS+
- IEEE 802.1x authentication and RADIUS
- Secure Shell (SSH) Protocol Version 2
- Simple Network Management Protocol Version 3 (SNMPv3)
- Port security
- DHCP snooping

Send comments to nx5000-docfeedback@cisco.com

- MAC ACLs and IP ACLs, including port-based ACLs (PACLs) and VLAN-based ACLs (VACLs).

Virtual Device Contexts

Cisco NX-OS can segment operating system and hardware resources into virtual device contexts (VDC) that emulate virtual devices. The Cisco Nexus 5000 Series switch does not support multiple VDCs. All switch resources are managed in the default VDC.

Typical Deployment Topologies

In this release, the Nexus 5020 switch is typically deployed in the following topologies:

- [Ethernet TOR Switch Topology, page 1-8](#)
- [IOC Topology, page 1-10](#)

Ethernet TOR Switch Topology


The Nexus 5020 switch can be deployed as a 10-Gigabit Ethernet top-of-rack (TOR) switch, with uplinks to the data center LAN distribution layer switches. An example configuration is shown in [Figure 1-2](#).

In this example, the blade server rack incorporates blade switches that support 10-Gigabit Ethernet uplinks to the Nexus 5020 switch. The blade switches do not support FCoE, so there is no FCoE traffic and no Fibre Channel ports on the Nexus 5020 switch.

In the example configuration, the Nexus 5020 switch has Ethernet uplinks to two Catalyst switches. If STP is enabled in the data center LAN, the links to one of the switches will be STP active and the links to the other switch will be STP blocked.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Figure 1-2 Ethernet TOR Switch Topology



All of the server-side ports on the Nexus 5020 switch are running standard Ethernet. FCoE is not required, so the server ports are connected using 10-Gigabit Ethernet NICs.

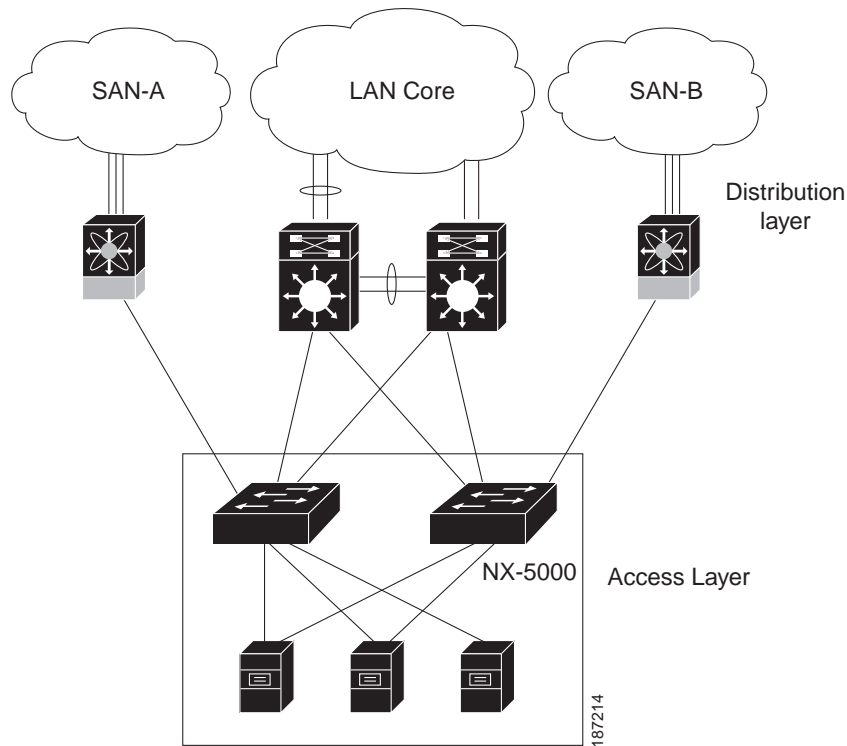
The servers are connected to the data center SAN through MDS 9134 SAN switches. The server Fibre Channel ports require standard Fibre Channel HBAs.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

IOC Topology

Figure 1-3 shows a typical I/O consolidation (IOC) scenario for the Nexus 5020 switch.

Figure 1-3 I/O Consolidation Topology



The Nexus 5020 switch connects to the server ports using FCoE. Ports on the server require converged network adapters. For redundancy, each server connects to both Nexus 5020 switches. Dual-port CNA adapters can be used for this purpose. The CNA is configured in active-passive mode, and the server needs to support server-based failover.

On the Nexus 5020 switch, the Ethernet network-facing ports are connected to two Catalyst 6500 switches. Depending on required uplink traffic volume, there may be multiple ports connected to each Catalyst 6500 switch, configured as port channels. If STP is enabled in the data center LAN, the links to one of the switches will be STP active and the links to the other switch will be STP blocked.

The Nexus 5020 SAN network-facing ports are connected to Cisco MDS 9000 Family switches. Depending on required traffic volume, there may be multiple Fibre Channel ports connected to each MDS 9000 Family switch, configured as SAN port channels.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Supported Standards

Table 1-1 lists the standards supported by the Cisco Nexus 5000 Series switches.

Table 1-1 *IEEE Compliance*

Standard	Description
802.1D	MAC Bridges
802.1s	Multiple Spanning Tree Protocol
802.1w	Rapid Spanning Tree Protocol
802.1AE	MAC Security (link layer cryptography)
802.3ad	Link aggregation with LACP
802.3ae	10 Gigabit Ethernet
802.1Q	VLAN Tagging
802.1p	Class of Service Tagging for Ethernet frames
802.1x	Port-based network access control

Send comments to nx5000-docfeedback@cisco.com



CHAPTER 2

Installing Cisco Fabric Manager

Cisco Fabric Manager provides an alternative to the command-line interface (CLI) for most switch configuration tasks. Fabric Manager provides a graphical user interface (GUI) that displays a real-time views of your Fibre Channel fabrics, and lets you manage the configuration of Cisco and third-party SAN switches.

This chapter contains the following sections:

- [Information About Cisco Fabric Manager, page 2-1](#)
- [Understanding Switch Management, page 2-3](#)
- [Installing the Management Software, page 2-5](#)
- [Upgrading the Management Software, page 2-18](#)
- [Integrating Cisco Fabric Manager with Other Management Tools, page 2-18](#)
- [Running Fabric Manager Behind a Firewall, page 2-19](#)
- [Uninstalling the Management Software, page 2-20](#)

Information About Cisco Fabric Manager

Cisco Fabric Manager provides an alternative to the command-line interface (CLI) for many switch configuration commands. For information about using the CLI to configure Cisco Nexus 5000 Series switches, see the *Cisco Nexus 5000 Series CLI Configuration Guide*.

In addition to complete configuration and status monitoring capabilities, Fabric Manager provides powerful Fibre Channel troubleshooting tools such as Fibre Channel ping and traceroute.

The Cisco Fabric Manager includes these management applications:

- [Fabric Manager Server, page 2-2](#)
- [Fabric Manager Client, page 2-2](#)
- [Fabric Manager Server Proxy Services, page 2-2](#)
- [Device Manager, page 2-3](#)
- [Performance Manager, page 2-3](#)
- [Fabric Manager Web Server, page 2-3](#)

Send comments to nx5000-docfeedback@cisco.com

Fabric Manager Server

The Fabric Manager Server software must be installed before running Fabric Manager. On a Windows PC, the Fabric Manager Server is installed as a service and is administered using the Windows Services control panel. Fabric Manager Server is responsible for discovery of the physical and logical fabric, and for listening for SNMP traps, syslog messages, and Performance Manager threshold events. For more information, see [Chapter 3, “Fabric Manager Server.”](#)

Fabric Manager Client

The Fabric Manager Client component displays a map of your Fibre Channel network fabrics, including Cisco switches, third-party switches, hosts, and storage devices. Fabric Manager Client provides multiple menus for accessing the features of Fabric Manager Server. For more information, see [Chapter 5, “Fabric Manager Client.”](#)

Fabric Manager Server Proxy Services

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

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

- External and internal loopback test
- Access to Flash files
- Creating CLI users
- **Show image version** command
- **Show tech** command
- Switch resident reports (syslog, accounting)
- Zone migration
- **Show cores** command

If you are using the SNMP proxy service and another application on your server is using port 9198, you need to modify your workstation settings.



Note

The switch always checks the local SNMP users before remote AAA users, unlike the CLI.

To modify a Windows workstation to use the proxy service, perform this task:

-
- Step 1** Open Internet Explorer and choose **Tools > Internet Options**.
You see the Internet Options dialog box.
- Step 2** Click the **Connections** tab and then choose **LAN Settings**.

Send comments to nx5000-docfeedback@cisco.com

You see the LAN Settings dialog box.

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

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

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

For additional information, see the [“Running Fabric Manager Behind a Firewall”](#) section on page 2-19.

Device Manager

Device Manager provides two views of the switch:

- Device View displays a graphic representation of the switch configuration and provides access to statistics and configuration information.
- Summary View displays a summary of xE ports (Inter-Switch Links), Fx ports (fabric ports), Nx ports (attached hosts and storage) and SAN port channels on the switch, as well as Fibre Channel and IP neighbor devices. Summary View also displays all operational virtual interfaces (Fibre Channel and Ethernet). Summary or detailed statistics can be charted, printed, or saved to a file in tab-delimited format. For additional information, see [Chapter 6, “Device Manager.”](#)

Performance Manager

Performance Manager provides detailed traffic analysis by capturing data with SNMP. This data is compiled into various graphs and charts that can be viewed with any web browser. For additional information, see [Chapter 29, “Performance Manager.”](#)

Fabric Manager Web Server

Fabric Manager Web Server allows operators to monitor and obtain reports for events, performance, and inventory from a remote location using a web browser.

Understanding Switch Management

Cisco Cisco Nexus 5000 Series switches are accessed and configured using standard management protocols. [Table 2-1](#) lists the management protocols that Fabric Manager supports to access, monitor, and configure Cisco Cisco Nexus 5000 Series switches.

Table 2-1 *Supported Management Protocols*

Management Protocol	Purpose
Telnet/SSH	Provides remote access to the CLI for a Cisco Cisco Nexus 5000 Series switch.
FTP/SFTP/TFTP, SCP	Copies configuration and software images between devices.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Table 2-1 Supported Management Protocols

Management Protocol	Purpose
SNMPv1, v2c, and v3	Includes over 80 distinct Management Information Bases (MIBs). Cisco Cisco Nexus 5000 Series 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 operation. By default, the Cisco Fabric Manager communicates with Cisco Cisco Nexus 5000 Series switches using SNMPv3, which provides secure authentication using encrypted user names and passwords. SNMPv3 also provides the option to encrypt all management traffic.
HTTP/HTTPS	Includes HTTP and HTTPS for web browsers to communicate with Fabric Manager Web Services and for the distribution and installation of the Cisco Fabric Manager software. HTTP is not used for communication between the Cisco Fabric Manager Server and Cisco Cisco Nexus 5000 Series switches.
XML/CIM over HTTP/HTTPS	Includes CIM server support for designing storage area network management applications to run on Cisco NX-OS.
ANSI T11 FC-GS-3	Provides Fibre Channel-Generic Services (FC-GS-3) in the defining management servers in the Fabric Configuration Server (FCS). Fabric Manager uses the information provided by FCS on top of the information contained in the Name Server database and in the Fibre Channel Shortest Path First (FSPF) topology database to build a detailed topology view and collect information for all the devices in the fabric.

Fabric Manager connects to switches using in-band or out-of-band management connections. These connection methods are described in the following sections:

- [mgmt0, page 2-4](#)
- [IPFC, page 2-5](#)

mgmt0

The out-of-band management connection is a 10/100 Mbps Ethernet interface on the supervisor module, which is 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 or Cisco

Send comments to nx5000-docfeedback@cisco.com

Nexus 5000 Series 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.

IPFC

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

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



Note

Fabric Manager requires an out-of-band (Ethernet) connection to at least one Cisco MDS 9000 Family or Cisco Nexus 5000 Series switch.

Installing the Management Software

This section describes how to install Fabric Manager and Device Manager. This section contains the following topics:

- [Before You Install, page 2-5](#)
- [Supported Software, page 2-6](#)
- [Installing the Database, page 2-7](#)
- [Installing Fabric Manager, page 2-8](#)
- [Installing Device Manager, page 2-16](#)

Before You Install

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

- Step 1** Configure the supervisor module with the following values using the setup routine or the CLI:
- IP address assigned to the mgmt0 interface
 - SNMP credentials (v3 user name and password or v1/v2 communities), maintaining the same user name and password for all the switches in the fabric



Note

Cisco Nexus 5000 Series switches support AAA authentication using RADIUS, TACACS, or local SNMP users.

- Step 2** Obtain the Fabric Manager software, which is available as a download from Cisco.com at the following location:

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

- Step 3** Shut down any instances of Fabric Manager and Device Manager.

Send comments to nx5000-docfeedback@cisco.com

**Note**

We recommend that you install the latest version of the Fabric Manager applications. Fabric Manager is backward compatible with the NX-OS software running on the Cisco Nexus 5000 Series switches. When upgrading the switch, upgrade the Fabric Manager software first, and then upgrade the Cisco NX-OS software on the switch.

Supported Software

**Note**

For the latest information on supported software, refer to the *Cisco Cisco Nexus 5000 Series Release Notes for Cisco NX-OS Release 4.0*.

Cisco Fabric Manager supports the following software:

- Operating systems
 - Windows 2000 SP4, 2003 SP2, XP SP2
 - Redhat Linux (2.6 Kernel)
 - Solaris (SPARC) 8 and 10
 - VMWare Server 1.0
- Java
 - Sun JRE and JDK 1.5(x) is supported
 - Java Web Start 1.2, 1.0.1 and 1.5
- Browsers
 - Internet Explorer 6.x and 7.0



Note Internet Explorer 7.0 is not supported on Windows 2000 SP4.

- Firefox 1.5 and 2.0
- Databases
 - Oracle Database 10g Express
 - PostgreSQL 8.2 (Windows)
 - PostgreSQL 8.1 (Solaris and Linux)
- Security
 - Cisco ACS 3.1 and 4.0
 - PIX firewall
 - IP tables
 - SSH v2
 - Global Enforce SNMP Privacy Encryption
 - HTTPS

Send comments to nx5000-docfeedback@cisco.com

Minimum Hardware Requirements

For a PC running Fabric Manager Server on large fabrics (1000 or more end devices), we recommend you use a Dual Core/Dual CPU high-speed system with 2 GB of RAM and 10 GB of free disk space.

Installing the Database

Fabric Manager requires an Oracle Database 10g Express or PostgreSQL database.

The Fabric Manager installation wizard provides an option to automatically install a PostgreSQL database. If you will not be selecting this option, you must install the database before you install Fabric Manager.



Note

We recommend the Oracle Database 10g Express option for all users who are running Performance Manager on large fabrics (1000 or more end devices).

Installing Oracle



Note

If you want to use Oracle Database 10g Express, you must install the database and create a user name and password before continuing with the Fabric Manager installation.

To install the Oracle database, perform this task:

Step 1 Go to the following location to install Oracle Database 10g Express:

<http://www.oracle.com/technology/software/products/database/xe/index.html>



Note

If you have another instance of Oracle already installed on a PC, we recommend that you do not install the Oracle database on the same PC. In such cases, Fabric Manager can only use the PostgreSQL database.

Step 2 Run OracleXE.exe to install the Oracle database, and then set the password for the system user.

The database administrator uses the password to manage and administer Oracle Database 10g Express server, which is installed by the Oracle installer.

Step 3 Finish the installation and verify that both services (OracleServiceXE and OracleXETNSListener) are running from the Services window.

Step 4 Run the following script to change the default Oracle admin port and to create a database account:

```
C:\> cd c:\oracle\app\oracle\product\10.2.0\server\bin
C:\oracle\app\oracle\product\10.2.0\server\bin>sqlplus / as sysdba
SQL> exec dbms_xdb.sethttpport(8082);
SQL> GRANT CONNECT,RESOURCE,UNLIMITED TABLESPACE TO SCOTT IDENTIFIED BY
TIGER;
SQL> EXIT;
```

Send comments to nx5000-docfeedback@cisco.com



Note The Oracle Database 10g Express option is only supported on Microsoft Windows. It is not supported on UNIX systems.

For information about backing up the Oracle database, go to the following location:

http://download.oracle.com/docs/cd/B25329_01/doc/admin.102/b25107/backrest.htm#i1004902

or use the exp/imp utility at:

http://download.oracle.com/docs/cd/B25329_01/doc/admin.102/b25107/impexp.htm#BCEEDCIB.



Note For information about backing up the PostgreSQL database, run the pg_dump utility. To run the utility, go to the following location: <http://www.postgresql.org/docs/8.1/static/app-pgdump.html>.

Installing Fabric Manager

You must install Fabric Manager from the CD-ROM or from Cisco.com.



Note Users installing Fabric Manager must have full administrator privileges to create user accounts and start services. Users should also have access to all ports. These are the ports used by Fabric Manager Server and the PostgreSQL database: 1098, 1099, 4444, 4445, 8009, 8083, 8090, 8092, 8093, 514, 5432.

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

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

To install Fabric Manager on Solaris, perform this task:

- Step 1 Set Java 1.5 to the path that is to be used for installing Fabric Manager.
- Step 2 Install the database that is to be used with Fabric Manager by following the instructions in the “[Installing the Database](#)” section on page 2-7.
- Step 3 Copy the Fabric Manager jar file m9000-fm-3.3.0.xx.jar from the CD-ROM to a folder on the Solaris workstation.



Note The filename on the CD-ROM will contain a number where “xx” is shown in step 3. Use the same number for “xx” in the **java** command in step 4.

- Step 4 Launch the installer using the following command:

```
java -Xmx256m -jar m9000-fm-3.3.0.xx.jar
```
- Step 5 Follow the on-screen instructions provided in the Fabric Manager management software setup wizard.

When you connect to the server for the first time, Fabric Manager checks to see if you have the correct Sun Java Virtual Machine version installed on your workstation. Fabric Manager looks for version 1.5(x) during installation. If required, install the Sun Java Virtual Machine software.

Send comments to nx5000-docfeedback@cisco.com



Note

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



Note

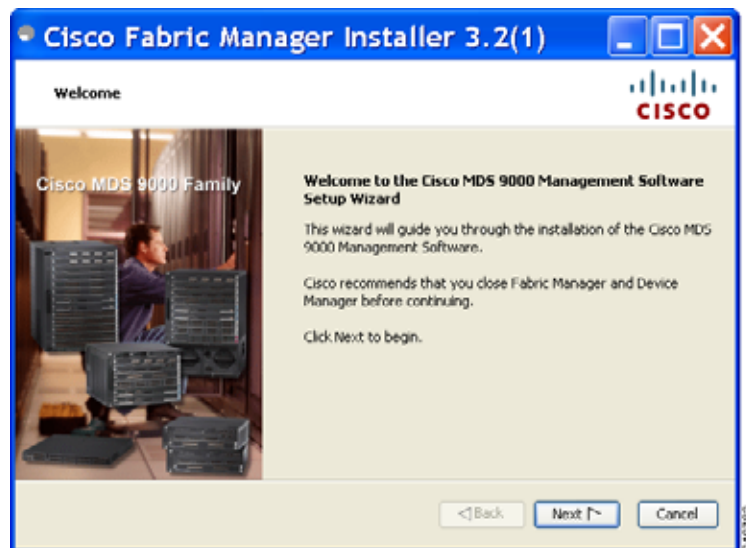
On Windows, Fabric Manager installations or upgrades should be done through the console using VNC and not a remote desktop.

To install Fabric Manager on Windows, perform this task:

- Step 1** Click the **Install Management Software** link.
- Step 2** Choose **Management Software > Cisco Fabric Manager**.
- Step 3** Click the **Installing Fabric Manager** link.
- Step 4** Click the **FM Installer** link.

You see the welcome to the management software setup wizard message in the Cisco Fabric Manager Installer window as shown in [Figure 2-1](#).

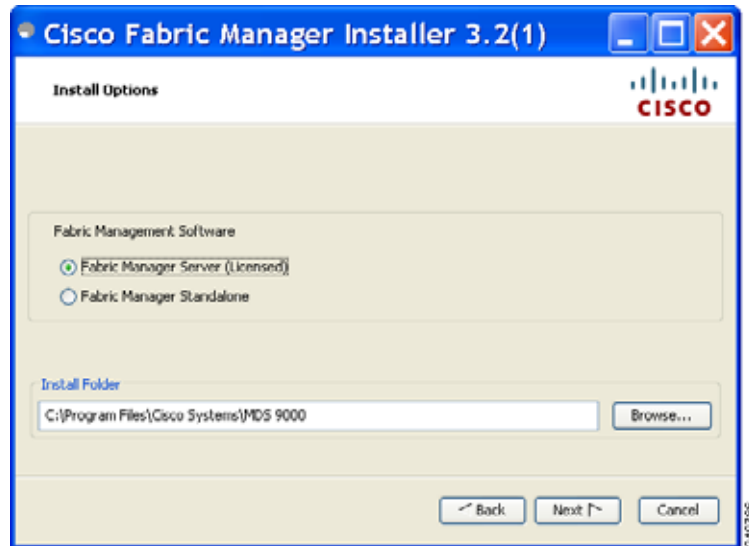
Figure 2-1 Welcome to the Management Software Setup Wizard



- Step 5** Click **Next** to begin the installation.
- Step 6** Check the **I accept the terms of the License Agreement** check box and click **Next**.
You see the Install Options dialog box as shown in [Figure 2-2](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 2-2 Install Options Dialog Box



Step 7 Click one of the radio buttons:

- Fabric Manager Server (Licensed) to install the server components for Fabric Manager Server.
- Fabric Manager Standalone to install the standalone version of Fabric Manager.



Note Fabric Manager Standalone is a single application containing Fabric Manager Client and a local version of Fabric Manager Server bundled together. Fabric Manager Standalone allows you to discover and monitor the immediate fabric.

Step 8 Select an installation folder on your workstation for Fabric Manager:

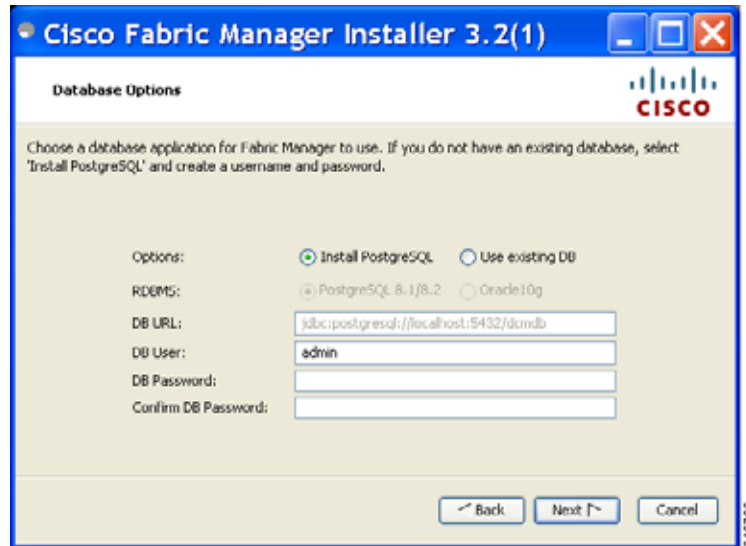
- On Windows, the default location is C:\Program Files\Cisco Systems\MDS 9000.
- On a UNIX (Solaris or Linux) machine, the installation path name is /usr/local/cisco_mds9000 or \$HOME/cisco_mds9000, depending on the permissions of the user doing the installation.

Step 9 Click **Next**.

You see the Database Options dialog box as shown in [Figure 2-3](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 2-3 Database Options Dialog Box



- Step 10** Click the **Install PostgreSQL** radio button to install Postgre SQL database. Click **Use existing DB** to specify which database you want to use.

If you choose Install PostgreSQL, accept the defaults and enter a password. The PostgreSQL database will be installed.



Note If you choose to install PostgreSQL, you must disable any security software you are running, because PostgreSQL may not install certain folders or users.

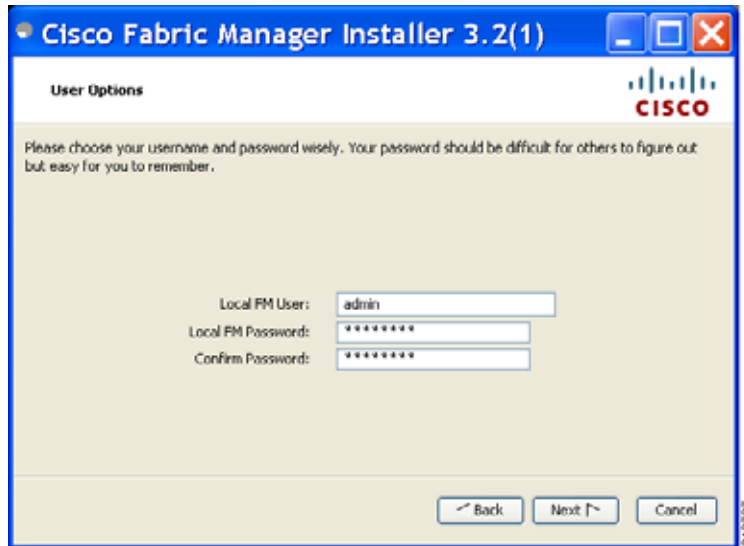


Note Before you install PostgreSQL, remove the **cygwin/bin** from your environment variable path if Cygwin is running on your system.

- Step 11** If you choose **Use existing DB**, click the **PostgreSQL 8.1/8.2** radio button or the **Oracle10g** radio button.
- Step 12** Click **Next** in the Database Options dialog box (figure [Figure 2-3](#)). You see the User Options dialog box as shown in [Figure 2-4](#).

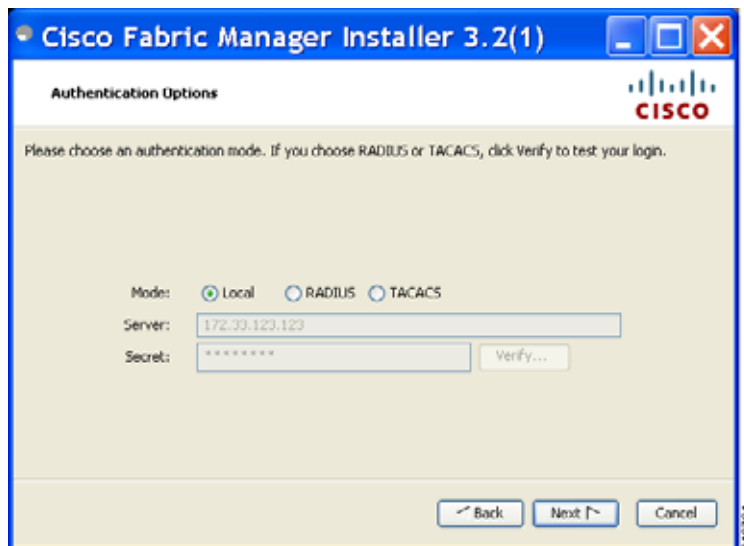
Send comments to nx5000-docfeedback@cisco.com

Figure 2-4 User Options Dialog Box



- Step 13** Enter a user name and password and click **Next**.
You see the Authentication Options dialog box as shown in [Figure 2-5](#).

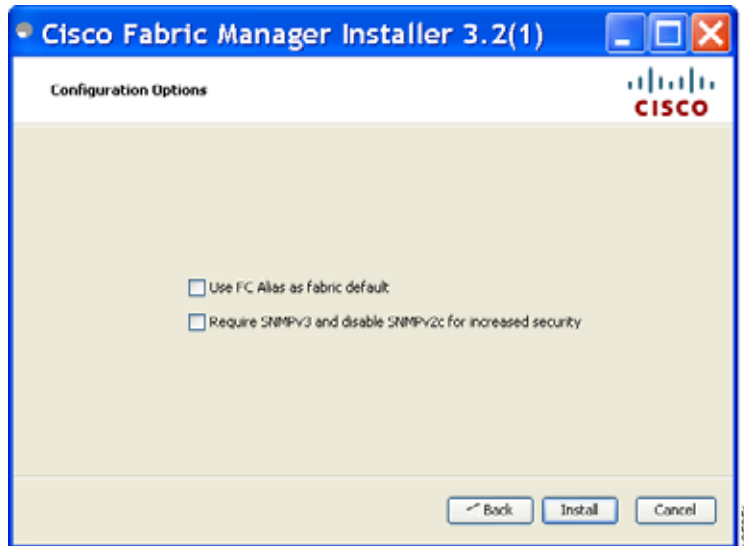
Figure 2-5 Authentication Options Dialog Box



- Step 14** Choose an authentication mode (Local, RADIUS, or TACACS) and click **Next**. Click **Verify** to test your login.
You see the Configuration Options dialog box for Fabric Manager Standalone as shown in [Figure 2-6](#).

Send comments to nx5000-docfeedback@cisco.com

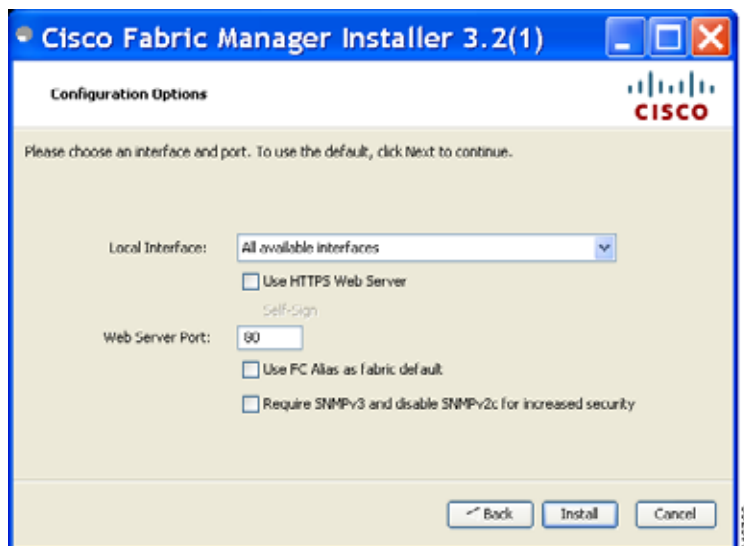
Figure 2-6 Configuration Options Dialog Box for Fabric Manager Standalone



- Step 15** Check the **FC Alias** and **SNMPv3** check boxes as desired and click **Install** if you are installing Fabric Manager Standalone.

You see the Configuration Options dialog box for Fabric Manager Server as shown in [Figure 2-7](#).

Figure 2-7 Configuration Options Dialog Box for Fabric Manager Server



- Step 16** Choose the local interface and web server port or check the **FC Alias** and **SNMPv3** check boxes as desired and click **Install** if you are installing Fabric Manager Server.



Note If you check the **Use HTTPS Web Server** check box, the Web Server Port field is grayed out and the default port is 443.

Send comments to nx5000-docfeedback@cisco.com

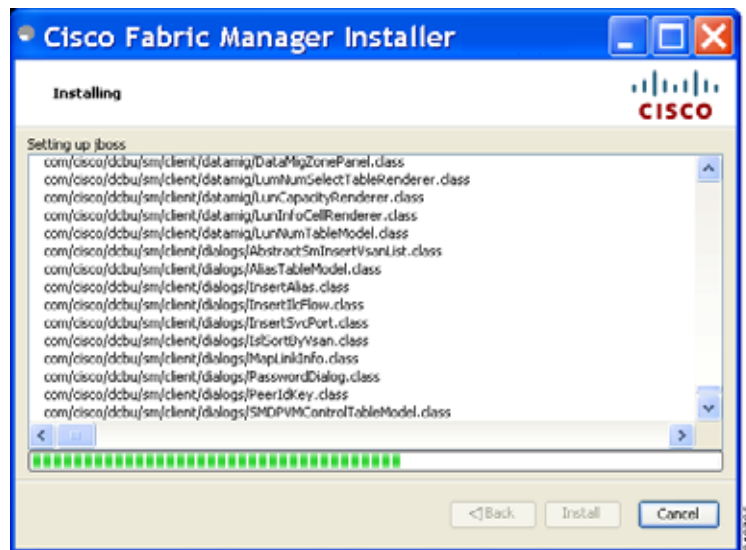


Note If you select a specific IP address during installation and change the server host IP address, you must modify the following two files, which are all located in the \$INSTALL/conf directory: Change **server.bindaddr**s to the new IP address in the server.properties file and change **wrapper.app.parameter.4** to the new IP address in the FMServer.conf file.

Step 17 Click **Cancel** to stop the installation.

You see the installation progress in the Cisco Fabric Manager Installer window as shown in [Figure 2-8](#).

Figure 2-8 Installation Progress



Once the installation is finished, you see an installation completed message in the Cisco Fabric Manager Installer window as shown in [Figure 2-9](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 2-9 Installation Complete



Note If you installed Fabric Manager Standalone, you can choose to launch Fabric Manager or Device Manager by checking the **Launch Fabric Manager** or **Launch Device Manager** check boxes. Icons for Fabric Manager and Device Manager are automatically created on the desktop.

Step 18 Click **Finish** to close the Cisco Fabric Manager Installer window.

If you installed Fabric Manager Server, icons for Fabric Manager and Device Manager are not created on the desktop until you launch Fabric Manager Client. Follow the instructions in the [“Launching Fabric Manager Client”](#) section on page 5-2 to launch Fabric Manager Client.

If you checked the Create shortcuts check box, a Cisco MDS 9000 program group is created under **Start > Programs** on Windows. This program group contains shortcuts to batch files in the install directory. Three services are started: Fabric Manager Server, Database, and Web Server. The Performance Manager server is installed but the service is not started upon installation because certain setup steps must be completed first.

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

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

Send comments to nx5000-docfeedback@cisco.com

Installing Device Manager

The Device Manager software executable file resides on the switch supervisor module. To install or upgrade the Device Manager software, access the supervisor module with a web browser and click the **Install** link on the web page that is displayed. The software running on your workstation is verified to make sure you are running the most current version of Device Manager. If it is not current, the most recent version is downloaded and installed on your workstation.

To install Device Manager on your workstation, perform this task:

-
- Step 1** Enter the IP address of the switch in the Address field of your browser.
You see the Installation window for Device Manager as shown in [Figure 2-10](#).

Figure 2-10 Device Manager Installation Window



- Step 2** Click the **Cisco Device Manager** link.
You see the welcome to the management software setup wizard message in the Cisco Device Manager Installer window as shown in [Figure 2-11](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 2-11 Welcome to the Management Software Setup Wizard Window



Step 3 Click **Next** to begin the installation.

Step 4 Check the **I accept the terms of the License Agreement** check box and click **Next**.


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

- On Windows, the default location is C:\Program Files\Cisco Systems\MDS 9000.
- On a UNIX (Solaris or Linux) machine, the installation path name is /usr/local/cisco_mds9000 or \$HOME/cisco_mds9000, depending on the permissions of the user doing the installation.

Step 6 Click **Install**.

You see the installation progress in the Cisco Device Manager Installer window as shown in [Figure 2-12](#).

Figure 2-12 Installation Progress



Send comments to nx5000-docfeedback@cisco.com

Once the installation is finished, you see an installation completed message in the Cisco Device Manager Installer window as shown in [Figure 2-13](#).

Figure 2-13 Installation Complete



Step 7 Click **Finish** to close the Cisco Device Manager Installer window.

Upgrading the Management Software

If you log into a switch with Device Manager and that switch has a later version of the Device Manager software, you are prompted to install the later version. You can also upgrade Device Manager at any time by entering the IP address or host name of the supervisor module with the later version of software in the Address field of your browser.



Note

Downgrades are not supported through the installer. To downgrade Device Manager to an earlier release, you need to manually uninstall first, and then reinstall the previous version of Device Manager.

To upgrade the Cisco Fabric Manager software, follow the instructions described in the [“Installing the Management Software”](#) section on [page 2-5](#). The installer supports software upgrade of the Fabric Manager application.

Integrating Cisco Fabric Manager with Other Management Tools

You can use Fabric Manager, Device Manager, and Performance Manager with these management tools:

- Cisco Traffic Analyzer—Allows you to break down traffic by VSANs and protocols and to examine SCSI traffic at a logical unit number (LUN) level.

Send comments to nx5000-docfeedback@cisco.com

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

For more information on these tools and how they work together with the Cisco Fabric Manager management applications, see [Chapter 31, “Troubleshooting Your Fabric.”](#)

Running Fabric Manager Behind a Firewall

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

By default, Fabric Manager Client and Device Manager use the first available UDP port for sending and receiving SNMP responses. The UDP SNMP trap local ports are 1162 for Fabric Manager, and 1163 or 1164 for Device Manager. Fabric Manager Server also opens TCP RMI port 9099.

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

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

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

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

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

Fabric Manager Server proxy services uses a configurable TCP port (9198 by default) for SNMP communications between the Fabric Manager Client or Device Manager and Fabric Manager Server.

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

- server.port = 9099
- server.data.port = 9100

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

[Table 2-2](#) lists all ports used by Fabric Manager applications.

Table 2-2 *Fabric Manager Port Usage*

Communication Type	Port(s) Used
Used by All Applications	
SSH	Port 22 (TCP)
Telnet	Port 23 (TCP)
HTTP	Port 80 (TCP)
TFTP	Port 69 (UDP)

Send comments to nx5000-docfeedback@cisco.com

Table 2-2 Fabric Manager Port Usage

Communication Type	Port(s) Used
Syslog	Port 514 (UDP)
Used by Fabric Manager Server and Performance Manager	
SNMP_TRAP	Port 2162 (UDP)
SNMP	Chooses a random free local port (UDP) or 9198 (TCP) if SNMP proxy is enabled. The port can be changed in server.properties.
Java RMI	Ports 9099, 9100 (TCP)
Used by Fabric Manager Client	
SNMP	Chooses a random free local port (UDP) if SNMP proxy is enabled. The port can be changed with the client -Dsnmp.localport option.
Java RMI	Chooses a free local port between 19199 and 19399 (TCP). The port can be changed with the client -Dclient.portStart and -Dclient.portEnd options. For example, -Dclient.portStart = 19199 -Dclient.portEnd = 19399.
Used by Device Manager	
SNMP_TRAP	Chooses a free local port between 1163 and 1170 (UDP).
SNMP	Chooses a random free local port (UDP) or 9198 (TCP) if SNMP proxy is enabled. The port can be changed in server.properties.

Uninstalling the Management Software

To uninstall the Fabric Manager applications on a Windows PC, perform this task:

-
- Step 1** Close all running instances of Fabric Manager and Device Manager.
- Step 2** Choose **Start > Programs > Cisco MDS 9000 > Uninstall** to run the uninstall.bat script.
- You can also run the batch file (located in the C:\Program Files\Cisco Systems\MDS 9000 folder by default) directly from the command line.
-

To uninstall the Fabric Manager applications on a UNIX machine, perform this task:

-
- Step 1** Run the shell script **\$HOME/cisco_mds9000/Uninstall.sh** or **/usr/local/cisco_mds9000/uninstall.sh**, depending on where Fabric Manager was installed.
-



- Note** Do not delete the MDS 9000 folder because this might prevent your installation from being upgraded in the future.
-



CHAPTER 3

Fabric Manager Server

Fabric Manager Server is a platform for advanced SAN monitoring, troubleshooting, and configuration capabilities. The server capabilities are an integral part of the Cisco Fabric Manager application.

This chapter contains the following sections:

- [Information About Fabric Manager Server, page 3-1](#)
- [Installing and Configuring Fabric Manager Server, page 3-2](#)
- [Managing a Fabric Manager Server Fabric, page 3-3](#)
- [Fabric Manager Server Properties File, page 3-4](#)
- [Modifying Fabric Manager Server, page 3-5](#)

Information About Fabric Manager Server

Install Cisco Fabric Manager Server on a computer that you want to provide centralized management services and performance monitoring. SNMP operations are used to efficiently collect fabric information. The Cisco Fabric Manager software, including the server components, requires about 60 MB of hard disk space on your workstation. Cisco Fabric Manager Server runs on Windows 2000, Windows 2003, Windows XP, Solaris 8.x or later, and Red Hat Linux.

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

Cisco Fabric Manager Server has the following features:

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



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

- **Continuous health monitoring**—Switch health is monitored continuously, so any events that occurred since the last time you opened the Fabric Manager Client are captured.

Send comments to nx5000-docfeedback@cisco.com

- Roaming user profiles—The licensed Fabric Manager Server uses the roaming user profile feature to store your preferences and topology map layouts on the server, so that your user interface will be consistent regardless of what computer you use to manage your storage networks.



Note

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

Installing and Configuring Fabric Manager Server

This section covers the installation and configuration of Fabric Manager server, and includes the following topics:

- [Installing Fabric Manager Server, page 3-2](#)
- [Unlicensed Versus Licensed Fabric Manager Server, page 3-3](#)
- [Installing Fabric Manager Web Server, page 3-3](#)
- [Verifying Performance Manager Collections, page 3-3](#)

Installing Fabric Manager Server

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

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



Note

Before running Fabric Manager Server, you should create a special Fabric Manager administrative user on each switch in the fabric or on a remote AAA server. Use this user to discover your fabric topology. See the [“Discovering a Fabric”](#) section on page 4-3.

To install Fabric Manager Server and set the initial configuration, perform this task:

- Step 1 Install Fabric Manager and Fabric Manager server on your workstation. See the [“Installing Fabric Manager Server”](#) section on page 3-2.
- Step 2 Log in to Fabric Manager. See the [“Launching Fabric Manager Client”](#) section on page 5-2.
- Step 3 Optionally, create flows Performance Manager to monitor your fabric. See the [“Counting Flow Statistics”](#) section on page 29-4.
- Step 4 Set Fabric Manager Server to continuously monitor the fabric. See the [“Managing a Fabric Manager Server Fabric”](#) section on page 3-3.
- Step 5 Repeat [Step 2](#) through [Step 4](#) for each fabric that you want to manage through Fabric Manager Server.
- Step 6 Install Fabric Manager Web Server. See the [“Installing Fabric Manager Web Server”](#) section on page 3-3.

Send comments to nx5000-docfeedback@cisco.com

- Step 7** Verify that Performance Manager is collecting data. See the “[Verifying Performance Manager Collections](#)” section on page 3-3.
-

Unlicensed Versus Licensed Fabric Manager Server

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

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

If you are evaluating one of these Fabric Manager Server features, use Device Manager to stop the evaluation period for that feature.

Installing Fabric Manager Web Server

You must install Fabric Manager Web Server to view Performance Manager reports through a web browser. To install Fabric Manager Web Server from the CD-ROM, see the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

Verifying Performance Manager Collections

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



Note

Viewing reports requires installing Fabric Manager Web Server. See the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

Managing a Fabric Manager Server Fabric

You can continuously manage a Fabric Manager Server fabric, whether or not a client has that fabric open. A continuously managed fabric is automatically reloaded and managed by Fabric Manager Server whenever the server starts.

Selecting a Fabric to Manage Continuously

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

Send comments to nx5000-docfeedback@cisco.com

To continuously manage a fabric using Fabric Manager, perform this task:

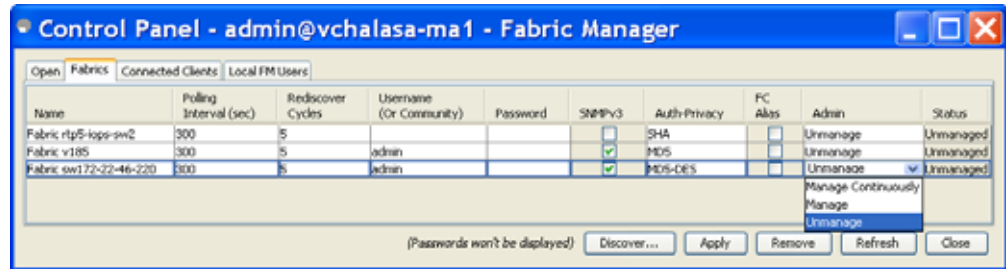
Step 1 Choose **Server > Admin**.

You see the Control Panel dialog box with the Fabrics tab open, as shown in [Figure 3-1](#).



Note The Fabrics tab is only accessible to network administrators.

Figure 3-1 Fabrics Tab in Control Panel Dialog Box



Note You can preconfigure a user name and password to manage fabrics. In this instance, you should use a local switch account, not a TACACS+ server.

Step 2 Select one of the following Admin options:

- **Manage Continuously**—The fabric is automatically managed when Fabric Manager Server starts and continues to be managed until this option is changed to Unmanage.
- **Manage**—The fabric is managed by Fabric Manager Server until there are no instances of Fabric Manager viewing the fabric.
- **Unmanage**—Fabric Manager Server stops managing this fabric.

Step 3 Click **Apply**.



Note If you are collecting data on these fabrics using Performance Manager, you should now configure flows and define the data collections. These procedures are described in [Chapter 29, “Performance Manager.”](#)

Fabric Manager Server Properties File

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



Note You can optionally encrypt the password in the server.properties and the AAA.properties files.

Send comments to nx5000-docfeedback@cisco.com

The server properties file contains these nine general sections:

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

The following are new or changed server properties for Fabric Manager Release 3.4:

- **Display FCoE**—Allows the user to display tree nodes, menu items, toolbar buttons, and topology nodes/links related to Fibre Channel over Ethernet (FCoE). If the fabric contains Cisco Nexus 5000 Series switches, set the **displayFCoE** property to true.



Note After you set the Display FCoE property, a Fabric Server restart is required.

Modifying Fabric Manager Server

Fabric Manager allows you to modify certain Fabric Manager Server settings without stopping and starting the server. These settings include:

- [Adding or Removing Fabric Manager Server Users, page 3-5](#)
- [Changing the Fabric Manager Server User Name and Password, page 3-6](#)
- [Changing the Polling Period and Fabric Rediscovery Time, page 3-6](#)

Adding or Removing Fabric Manager Server Users

To add a Fabric Manager Server user or to change the password for an existing user using Fabric Manager, perform this task:

- Step 1** Click the **Local FM Users** tab in the Control Panel dialog box, as shown in [Figure 3-1](#). You see a list of Fabric Manager users.



Note Only network administrators can manage users.

- Step 2** Click **New** to add a user or click the user name and click **Edit** to change the password for an existing user. You see the FM User dialog box as shown in [Figure 3-2](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 3-2 FM User Dialog Box

- Step 3** Set the user name and password for the new user, and then click **Apply**.
-

To remove a Fabric Manager Server user using Fabric Manager, perform this task:

- Step 1** Click the **Local FM Users** tab in the Control Panel dialog box (see [Figure 3-1](#)).
You see a list of Fabric Manager users.
- Step 2** Click the user name you want to delete.
- Step 3** Click **Remove** to delete the user.
- Step 4** Click **Yes** to confirm the deletion or **No** to cancel it.
-

Changing the Fabric Manager Server User Name and Password

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

To change the user name or password used by Fabric Manager Server, perform this task:

- Step 1** Choose **Server > Admin**.
You see the Control Panel dialog box with the Fabrics tab open, as shown in [Figure 3-1](#).
- Step 2** Set the Name or Password for each fabric that you are monitoring with Fabric Manager Server.
- Step 3** Click **Apply** to save these changes.
-

Changing the Polling Period and Fabric Rediscovery Time

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

Send comments to nx5000-docfeedback@cisco.com

To change the polling period or full fabric rediscovery setting used by Fabric Manager Server using Fabric Manager, perform this task:

-
- Step 1** Choose **Server > Admin**.
- You see the Control Panel dialog box with the Fabrics tab open, as shown in [Figure 3-1](#).
- Step 2** For each fabric that you are monitoring with Fabric Manager Server, set the Polling Interval to determine how frequently Fabric Manager Server polls the fabric elements for status and statistics.
- Step 3** For each fabric that you are monitoring with Fabric Manager Server, set the Rediscover Cycles to determine how often Fabric Manager Server rediscovers the full fabric.
- Step 4** Click **Apply** to save these changes.
-

Using Device Aliases or FC Aliases

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

To change whether Fabric Manager uses FC aliases or global device aliases using Fabric Manager, perform this task:

-
- Step 1** Choose **Server > Admin**.
- You see the Control Panel dialog box with the Fabrics tab open, as shown in [Figure 3-1](#).
- Step 2** For each fabric that you are monitoring with Fabric Manager Server, check the **Device Alias** check box to use global device aliases, or uncheck to use FC aliases.
- Step 3** Click **Apply** to save these changes.
-

Saving Device Aliases to the Switch

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

To save global device aliases on a switch using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches**, > **End Devices** and then choose **Hosts** or **Storage**.
- You see the end devices in the Information pane.
- Step 2** For each device alias that you want the switch to recognize, highlight it, right-click the **Device Alias** icon for that device, and choose **Save Selected Device Aliases**.
-

Send comments to nx5000-docfeedback@cisco.com



CHAPTER 4

Authentication in Fabric Manager

Fabric Manager contains interdependent software components that communicate with the switches in your fabric. These components use varying methods to authenticate to other components and switches. This chapter describes these authentication steps and the recommended authentication configuration for your fabric and components.

This chapter contains the following sections:

- [Information About Fabric Manager Authentication, page 4-1](#)
- [Discovering a Fabric, page 4-3](#)
- [Authenticating Performance Manager, page 4-4](#)
- [Authenticating Fabric Manager Web Server, page 4-5](#)

Information About Fabric Manager Authentication

Fabric Manager contains multiple components that interact to manage a fabric.


This chapter includes the following sections:

- Fabric Manager Client
- Fabric Manager Server
- Performance Manager
- Interconnected fabric of Cisco SAN switches and storage devices
- AAA server (optional)

[Figure 4-1](#) shows an example configuration for these components.

Send comments to nx5000-docfeedback@cisco.com

Figure 4-1 Fabric Manager Authentication Example



130715

Administrators launch Fabric Manager Client and select the seed switch that is used to discover the fabric. The user name and password used are passed to Fabric Manager Server and are used to authenticate access to the seed switch. If this user name and password are not a recognized SNMP user name and password, either Fabric Manager Client or Fabric Manager Server opens a CLI session to the switch (SSH or Telnet) and retries the user name and password pair. If the user name and password are recognized by the switch in either the local switch authentication database or through a remote AAA server, then the switch creates a temporary SNMP user name that is used by Fabric Manager Client and server.

**Note**

You may encounter a delay in authentication if you use a remote AAA server to authenticate Fabric Manager or Device Manager.

**Note**

You must allow CLI sessions to pass through any firewall that exists between Fabric Manager Client and Fabric Manager Server. See the [“Running Fabric Manager Behind a Firewall”](#) section on page 2-19.

**Note**

We recommend that you use the same password for the SNMPv3 user name authentication and privacy passwords as well as the matching CLI user name and password.

Send comments to nx5000-docfeedback@cisco.com

Discovering a Fabric

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



Caution

If the Fabric Manager Server's CPU usage exceeds 50 percent, it is recommended that you switch to a higher CPU-class system. For more information on recommended hardware, see the [“Before You Install” section on page 2-5](#).

We recommend that you use the steps described in the following sections for discovering your network and setting up Performance Manager. This procedure ensures that Fabric Manager Server has a complete view of the fabric. Subsequent Fabric Manager Client sessions can filter this complete view based on the privileges of the client logging in. For example, if you have multiple VSANs in your fabric and you create users that are limited to a subset of these VSANs, you want to initiate a fabric discovery through Fabric Manager Server using a network administrator or network operator role so that Fabric Manager Server has a view of all the VSANs in the fabric. When a VSAN-limited user launches Fabric Manager Client, that user sees only the VSANs they are allowed to manage.



Note

Fabric Manager Server should always monitor fabrics using a local switch account. Do not use a AAA (RADIUS or TACACS+) server. You can use a AAA user account to log into the clients to provision fabric services. For more information on Fabric Manager Server fabric monitoring, see the [“Managing a Fabric Manager Server Fabric” section on page 3-3](#).

Setting Up Discovery for a Fabric

To ensure that Fabric Manager Server discovers your complete fabric, perform this task:

-
- Step 1** Create a special Fabric Manager administrative user name in each switch on your fabric with network administrator or network operator roles.
You can alternatively create a special Fabric Manager administrative user name in your AAA server and set every switch in your fabric to use this AAA server for authentication.
 - Step 2** Verify that the roles used by this Fabric Manager administrative user name are the same on all switches in the fabric and that this role has access to all VSANs.
 - Step 3** Launch Fabric Manager Client using the Fabric Manager administrative user.
This step ensures that your fabric discovery includes all VSANs.
 - Step 4** Set Fabric Manager Server to continuously monitor the fabric.
See the [“Managing a Fabric Manager Server Fabric” section on page 3-3](#).
 - Step 5** Repeat [Step 4](#) for each fabric that you want to manage through Fabric Manager Server.
-

Send comments to nx5000-docfeedback@cisco.com

Authenticating Performance Manager

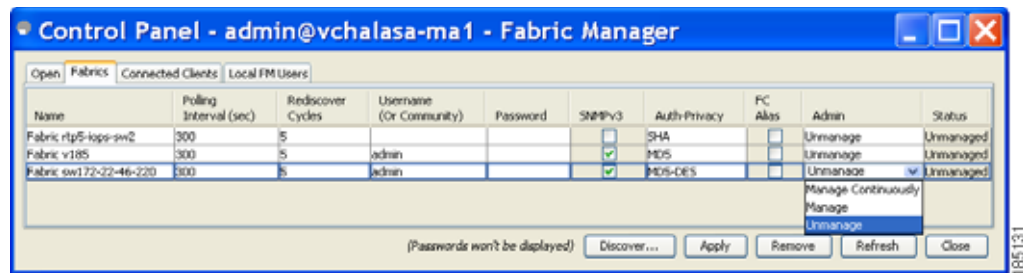
Performance Manager uses the user name and password information stored in the Fabric Manager Server database. If this information changes on the switches in your fabric while Performance Manager is running, you need to update the Fabric Manager Server database and restart Performance Manager. Updating the Fabric Manager Server database requires removing the fabric from Fabric Manager Server and rediscovering the fabric.

To update the user name and password information used by Performance Manager, perform this task:

Step 1 Click **Server > Admin** in Fabric Manager.

You see the Control Panel dialog box with the Fabrics tab open, as shown in [Figure 4-2](#).

Figure 4-2 Fabrics Tab in Control Panel Dialog Box



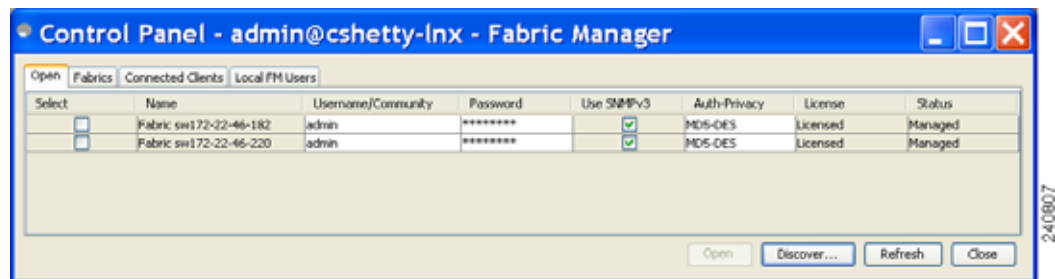
Step 2 Click the fabrics that have updated user name and password information.

Step 3 Click **Remove** to remove these fabrics from Fabric Manager Server.

Step 4 Choose **File > Open Fabric**.

You see the Control Panel dialog box as shown in [Figure 4-3](#).

Figure 4-3 Control Panel Dialog Box



Step 5 Enter the appropriate user name and password to rediscover the fabric and check the check boxes in the Select column next to the fabrics that you want to open.

Step 6 Click **Open** to rediscover the fabric.

Fabric Manager Server updates its user name and password information.

Step 7 Repeat [Step 4](#) through [Step 6](#) for any fabric that you need to rediscover.

Send comments to nx5000-docfeedback@cisco.com

- Step 8** Choose **Performance > Collector > Restart** to restart Performance Manager and use the new user name and password.
-

Authenticating Fabric Manager Web Server

Fabric Manager Web Server does not communicate directly with any switches in the fabric. Fabric Manager Web Server uses its own user name and password combination that is either stored locally or stored remotely on an AAA server.

We recommend that you use a RADIUS or TACACS+ server to authenticate users in Fabric Manager Web Server.

To configure Fabric Manager Web Server to use RADIUS authentication, perform this task:

- Step 1** Launch Fabric Manager Web Server.
- Step 2** Choose **Admin tab > Web Users** to update the authentication used by Fabric Manager Web Server.
- Step 3** Click **AAA**.
- Step 4** Set the authentication mode attribute to **radius**.
- Step 5** Set the RADIUS server name, shared secret, authentication method, and ports used for up to three RADIUS servers.
- Step 6** Click **Modify** to save this information.
-

To configure Fabric Manager Web Server to use TACACS+ authentication, perform this task:

- Step 1** Launch Fabric Manager Web Server.
- Step 2** Choose **Admin > Web Users** to update the authentication used by Fabric Manager Web Server.
- Step 3** Click **AAA**.
- Step 4** Set the authenticationmode attribute to **tacacs**.
- Step 5** Set the TACACS+ server name, shared secret, authentication method, and port used for up to three TACACS+ servers.
- Step 6** Click **Modify** to save this information.
-

Send comments to nx5000-docfeedback@cisco.com



CHAPTER 5

Fabric Manager Client

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

This chapter contains the following sections:

- [Information About Fabric Manager Client, page 5-1](#)
- [Launching Fabric Manager Client, page 5-2](#)
- [Fabric Manager Client Quick Tour, page 5-6](#)
- [Setting Fabric Manager Preferences, page 5-18](#)
- [Network Fabric Discovery, page 5-19](#)
- [Modifying the Device Grouping, page 5-20](#)
- [Controlling Administrator Access with Users and Roles, page 5-21](#)
- [Using Fabric Manager Wizards, page 5-21](#)
- [Fabric Manager Troubleshooting Tools, page 5-22](#)

Information About Fabric Manager Client

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

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



Note

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

Fabric Manager Advanced Mode

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

Send comments to nx5000-docfeedback@cisco.com

Launching Fabric Manager Client



Note

Network administrators must initially launch Fabric Manager Client using Fabric Manager Web Server, as described below. Once an administrator has installed the Fabric Manager Client icon on your desktop, you can double click on the icon to launch the Fabric Manager Client.

To launch Fabric Manager Client, perform this task:

- Step 1** Open your browser and enter the IP address where you installed Fabric Manager Server, or enter localhost if you installed Fabric Manager Server on your local workstation.

You see the Fabric Manager Web Server Login dialog box shown in [Figure 5-1](#).

Figure 5-1 Fabric Manager Web Server Login Dialog Box



- Step 2** Enter your user name and password and click **Login**.

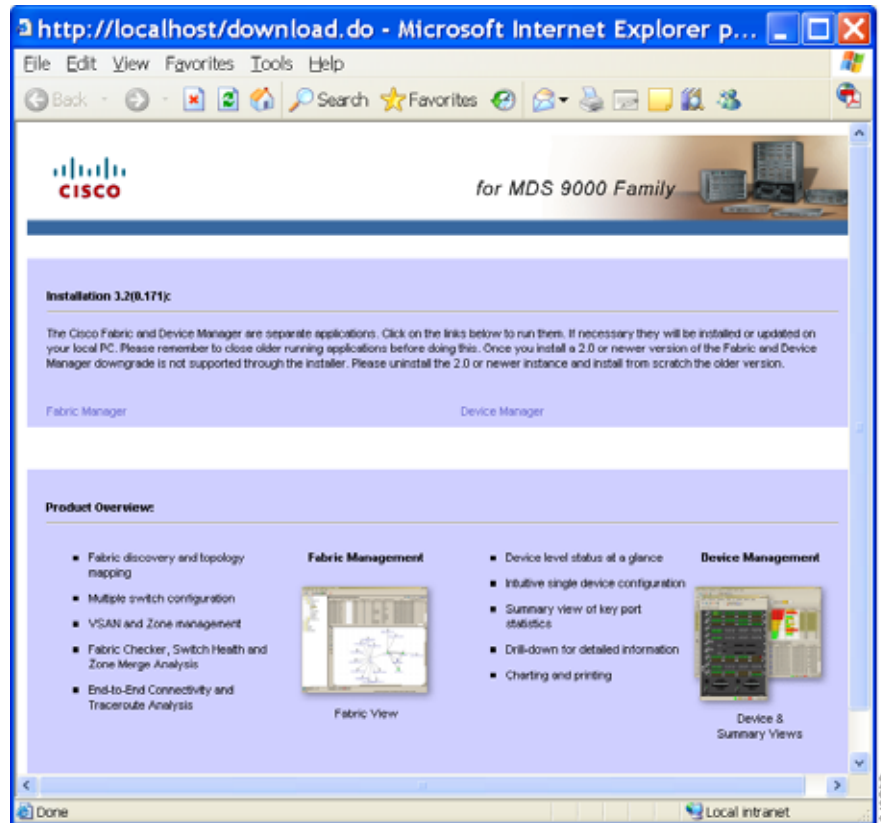
You see the Fabric Manager Web Server Summary page.

- Step 3** Click the **Download** link in the upper right corner of the page.

You see the Download page for Fabric Manager and Device Manager as shown in [Figure 5-2](#).

Send comments to nx5000-docfeedback@cisco.com


Figure 5-2 Download Page for Fabric Manager and Device Manager



Step 4 Click the link for either **Fabric Manager** or **Device Manager**.

If you are launching Fabric Manager Client for the first time, you see a message asking whether you want to create shortcuts for Fabric Manager, as shown in [Figure 5-3](#).

Figure 5-3 Fabric Manager Create Shortcut(s) Message



Step 5 Click **Yes** to create shortcuts for Fabric Manager.



Note This message only appears the first time you launch Fabric Manager Client.

Step 6 When the software is installed and icons are created on your desktop, double-click the Fabric Manager icon to launch Fabric Manager.

You see the Fabric Manager Login dialog box shown in [Figure 5-4](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 5-4 Fabric Manager Login Dialog Box



- Step 7** Enter the Fabric Manager Server user name and password.
- Step 8** Check the **Use SNMP Proxy** check box if you want Fabric Manager Client to communicate with Fabric Manager Server through a TCP-based proxy server.
- Step 9** Click **Login**. After you successfully log in to Fabric Manager Server, you can set the seed switch and open the fabrics that you are entitled to access.



Note When you launch Fabric Manager Client for the first time or when there are no available fabrics, you see the Discover New Fabric dialog box.

You see the Discover New Fabric dialog box shown in [Figure 5-5](#).

Figure 5-5 Discover New Fabric Dialog Box



Note Only network administrators can discover new fabrics.

- Step 10** Set the fabric seed switch to the Cisco switch that you want Fabric Manager to use.



Note A Cisco switch running in NPV mode cannot be the fabric seed switch.

- Step 11** Enter the user name and password for the switch.
- Step 12** Choose the Auth-Privacy option according to the privacy protocol you have configured on your switch:

Send comments to nx5000-docfeedback@cisco.com

- If you have not configured the switch with a privacy protocol, then choose Auth-Privacy option MD5 (no privacy).
- If you have configured the switch with your privacy protocol, choose your Auth-Privacy option.

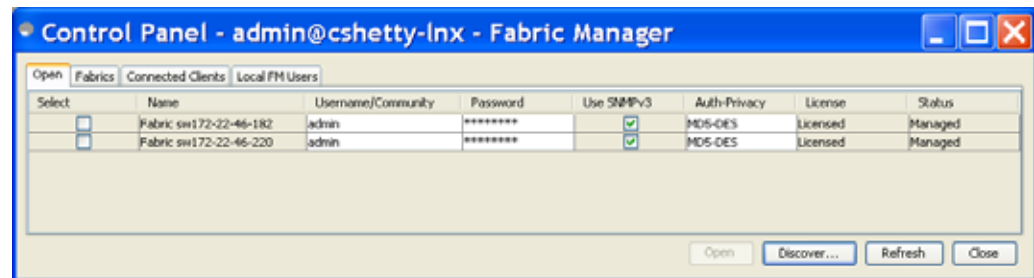


Note If you want a clean discovery, remove the fabric and rediscover it.

Step 13 Click **Discover**.

You see the Control Panel dialog box shown in [Figure 5-6](#).

Figure 5-6 Control Panel Dialog Box



Note You see a message in the dialog box when the server and client are running on the same workstation and there are unlicensed fabrics in the database. You also see a message when there are unmanaged fabrics (the state of the licenses is unknown).

Step 14 Check the check boxes in the Select column next to the fabrics that you want to open, or click **Discover** to add a new fabric.



Note Only network administrators can continuously manage or unmanage fabrics. For more information, see the [“Selecting a Fabric to Manage Continuously”](#) section on page 3-3.

Step 15 Click **Open** to open the selected fabrics.



Note If you have an incomplete view of your fabric, rediscover the fabric with a user that has no VSAN restriction.

To launch Fabric Manager Client from within a running instance of Fabric Manager, perform this task:

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

You see the Control Panel dialog box as shown in [Figure 5-6](#).

Step 2 Check the check boxes in the Select column next to the fabrics you want to open and click **Open**.

Send comments to nx5000-docfeedback@cisco.com

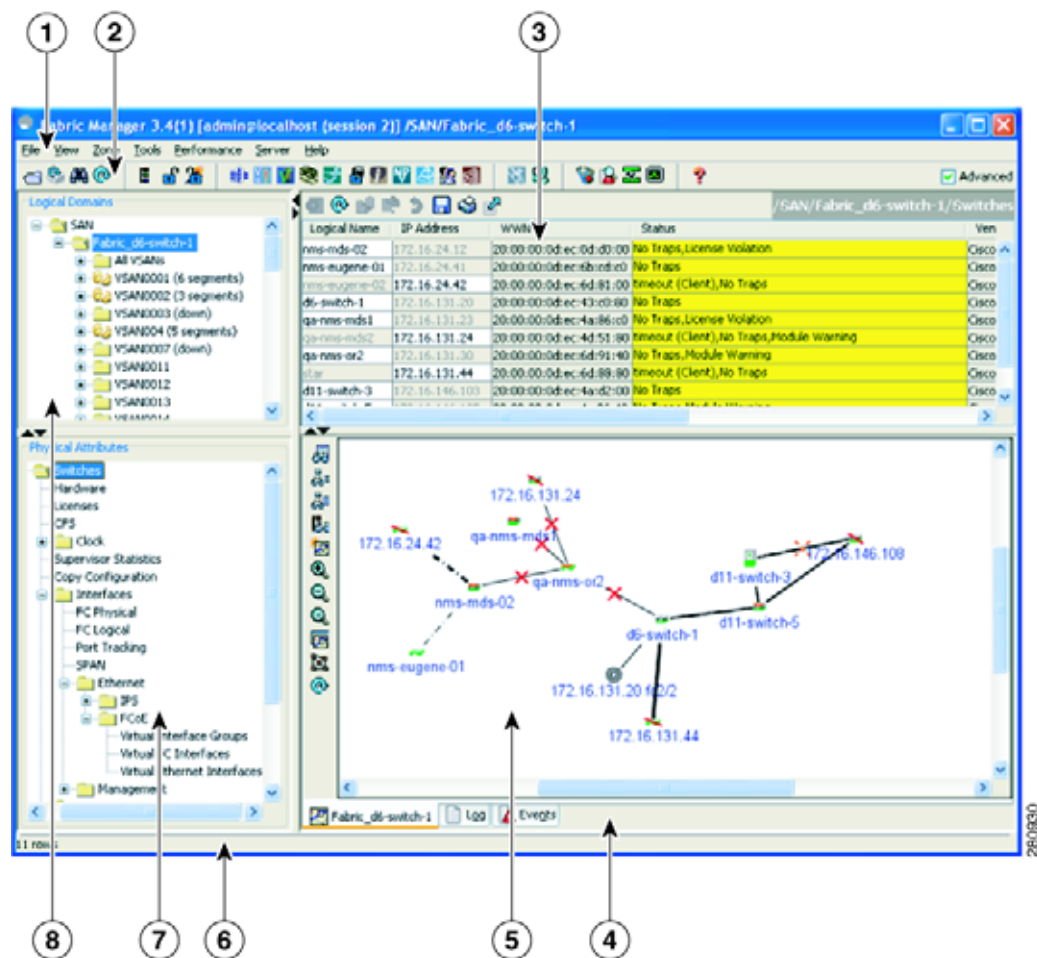
**Note**

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

Fabric Manager Client Quick Tour

This section describes the Fabric Manager Client interface shown in [Figure 5-7](#).

Figure 5-7 Fabric Manager Main Window



1	Menu bar—Provides access to options that are organized by menus.
2	Toolbar—Provides icons for direct access to the most commonly used options on the File, Tools, and Help menus.
3	Information pane—Displays information about whatever option is selected in the menu tree.

Send comments to nx5000-docfeedback@cisco.com

4	Status Bar (right side)—Shows the last entry displayed by the discovery process and the possible error message.
5	Fabric pane—Displays a map of the network fabric, including switches, hosts, and storage. It also provides tabs for displaying log and event data.
6	Status Bar (left side)—Shows short-term transient messages, such as the number of rows displayed in a table.
7	Logical Domains pane—Displays a tree of configured SAN, fabrics, VSANs, and zones, and provides access to user-defined groups.
8	Physical Attributes pane—Displays a tree of available configuration tasks depending on the fabric, VSAN, or zone selected previously. Lists the switches and end devices in the logical selection.



Note

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

Menu Bar

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




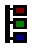








- **File**—Opens a new fabric, rediscovers the current fabric, locates switches, sets preferences, prints the map, and clears (right-click on log) or exports the Fabric pane log.
- **View**—Changes the appearance of the map (these options are duplicated on the Fabric pane toolbar).
- **Zone**—Manages zones, zone sets, and inter-VSAN routing (IVR).
- **Tools**—Verifies and troubleshoots connectivity and configuration, as described in the “[Fabric Manager Troubleshooting Tools](#)” section on page 5-22.
- **Performance**—Runs and configures Performance Manager and Cisco Traffic Analyzer and generates reports.
- **Server**—Runs administrative tasks on clients and fabrics. Provides Fabric Manager Server management and a **purge** command. Lists switches being managed.
- **Help**—Displays online help topics for specific dialog boxes in the Information pane.

Send comments to nx5000-docfeedback@cisco.com

Toolbar




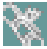






The Fabric Manager main toolbar provides icons for accessing the most commonly used menu bar options as shown in [Table 5-1](#).

Table 5-1 *Fabric Manager Client Main Toolbar*

Icon	Description
	Opens switch fabric.
	Rediscovered current fabric.
	Searches the map.
	Creates VSAN.
	Launches DPVM wizard.
	Edits full zone database.
	Launches IVR zone wizard.
	Launches SAN port channel wizard.
	Launches Virtual Interface Group wizard.
	Launches Virtual Interface wizard.
	Launches FCIP wizard.
	Launches iSCSI wizard.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Table 5-1 Fabric Manager Client Main Toolbar (continued)

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

Logical Domains Pane

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

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

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

Send comments to nx5000-docfeedback@cisco.com

To change the fabric name using Fabric Manager, perform this task:

Step 1 Choose **Server > Admin**.

You see the Control Panel dialog box.

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

Step 3 Click **Apply** to change the name.

Filtering

Fabric Manager has a filtering mechanism that displays only the data that you are interested in. To filter data, first select the fabric and VSAN from the Logical Domains pane. This action narrows the scope of what is displayed in the Fabric pane. Any information that does not belong to the selected items is dimmed. Also, any information that does not belong to the selected items is not displayed in the tables in the Information pane. The filter that you select is displayed at the top right of the Fabric Manager window (see [Figure 5-8](#)).

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

Physical Attributes Pane

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

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











- Switches—Views and configures hardware, system, licensing, and configuration files.
- Interfaces—Views and configures physical Fibre Channel, virtual Fibre Channel, Ethernet, SVC, and SAN port channel interfaces.
- Fibre Channel Services—Views and configures Fibre Channel network configurations.
- Events—Views and configures events, alarms, thresholds, notifications, and informs.
- Security—Views and configures switch management and FC-SP security.
- End Devices—Views and configures end devices.

Send comments to nx5000-docfeedback@cisco.com

Information Pane

Use the Information pane to display tables of information associated with the option selected from the menu tree in the Logical Domains or Physical Attributes panes. The Information pane toolbar provides buttons for performing one or more of the operations shown in [Table 5-2](#).

Table 5-2 Information Pane Toolbar

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



Note

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

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

**Note**

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









Detachable Tables

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

Fabric Pane














Use the Fabric pane to display the graphical representation of your fabric. [Table 5-3](#) describes the graphics you may see displayed, depending on which devices you have in your fabric.

Table 5-3 *Fabric Manager Graphics*

Icon or Graphic	Description
	Director class Cisco MDS 9000.
	Non-director class switch.
	Generic Fibre Channel switch.
	Cisco SN5428.
	Cisco Nexus 5000 series switch.
	An orange line through a device indicates that the device is manageable but there are operational problems.
	An orange X through a device or link indicates that the device or ISL is not working properly.
	A red line through a device indicates that the device is not manageable.

Send comments to nx5000-docfeedback@cisco.com

Table 5-3 *Fabric Manager Graphics (continued)*

Icon or Graphic	Description
	A red X through a device or link indicates that the device is down or that the ISL is down.
	Fibre Channel HBA (or enclosure).
	Fibre Channel target (or enclosure).
	iSCSI host.
	Fibre Channel ISL and edge connection.
	Fibre Channel SAN port channel.
	Fibre Channel over Ethernet (FCoE) connection.
	IP ISL and edge connection.
	IP port channel.
	NPV connection.
	Fibre Channel loop (storage).
	IP cloud (hosts). This icon is also used to represent a fabric when viewing a SAN (multiple fabrics) in the Fabric Manager Fabric pane.
	Any device, cloud, or loop with a box around it means that there are hidden links attached.

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

The bottom of the Fabric pane has the following tabs:

- Fabric—When displaying multiple fabrics, each fabric has its own tab. You can switch between fabrics by clicking on their respective tabs.

Send comments to nx5000-docfeedback@cisco.com

- Log—Displays messages that describe Fabric Manager operations, such as fabric discovery.
- Events—Displays information about the SNMP traps received by the management station. This includes combination events as detected by discovery and important traps, such as license and SNMP.

You can view large fabrics in the Fabric pane easier if you do the following tasks:

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



Note

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

Context Menus

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

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

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



Note

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

Saving the Map

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

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

Send comments to nx5000-docfeedback@cisco.com

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

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

Purging Down Elements

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

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

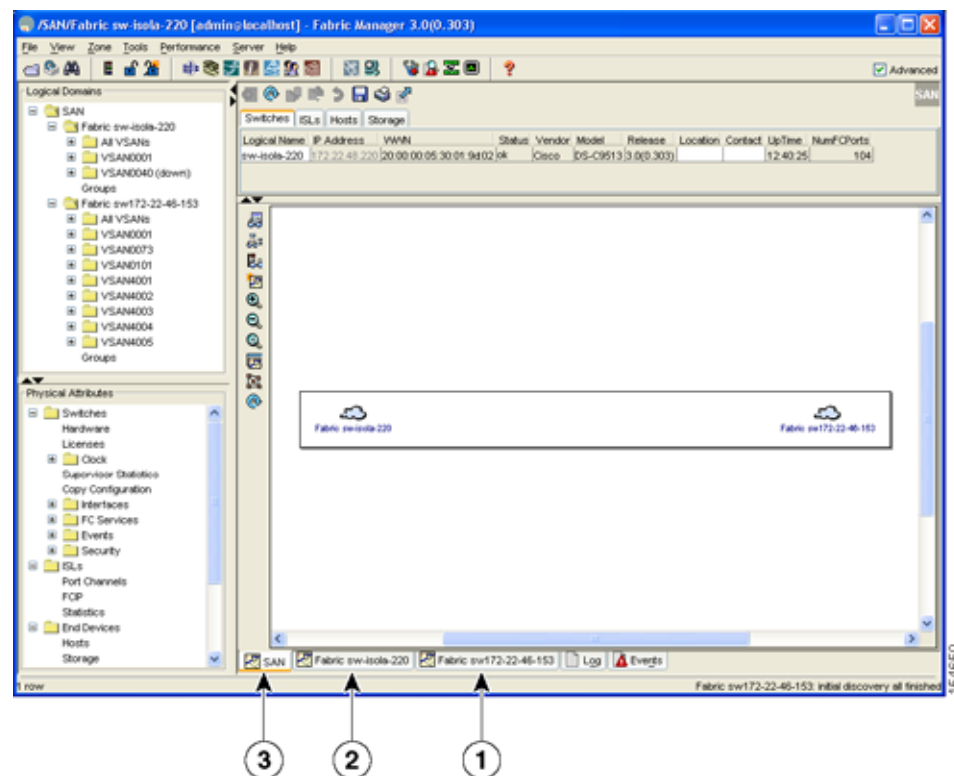


Note If you select an element that is not down and you purge it, that element will reappear on the next fabric discovery cycle.

Multiple Fabric Display

Fabric Manager can display multiple fabrics in the same pane (see [Figure 5-8](#)).

Figure 5-8 Fabric Manager's Multiple Fabric Display Window



Send comments to nx5000-docfeedback@cisco.com

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

**Note**

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

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

**Note**

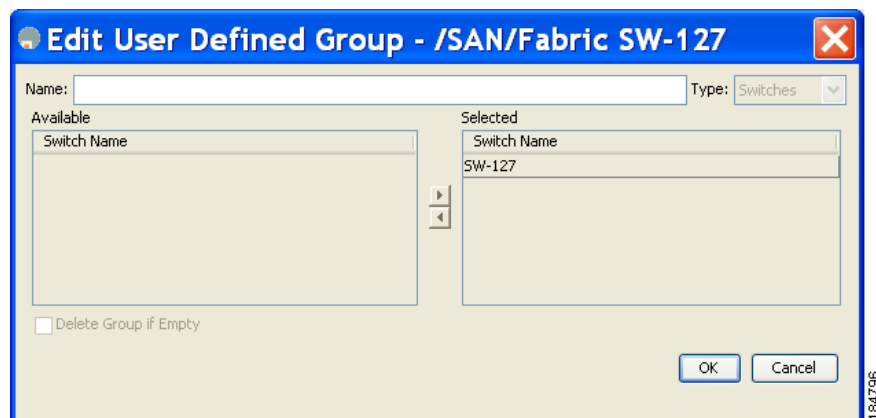
Enclosure names should be unique. If the same enclosure name is used for each port, Fabric Manager shows a host/target enclosure connected to both fabrics. To fix this problem, you can either disable auto-creation or create unique enclosure names.

Filtering by Groups

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

- Step 1** Right-click a switch or end port in the Fabric pane map and choose **Group > Create**. You see the Edit User Defined Group dialog box shown in [Figure 5-9](#).

Figure 5-9 Edit User Defined Group Dialog Box



- Step 2** Enter a group name in the **Name** field.
- Step 3** Use the arrows to move additional switches or end ports from the Available column to the Selected column.

Send comments to nx5000-docfeedback@cisco.com

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

To add a switch or end port to an existing group in Fabric Manager, perform this task:

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

You see the Edit User Defined Group dialog box as shown in [Figure 5-9](#).

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

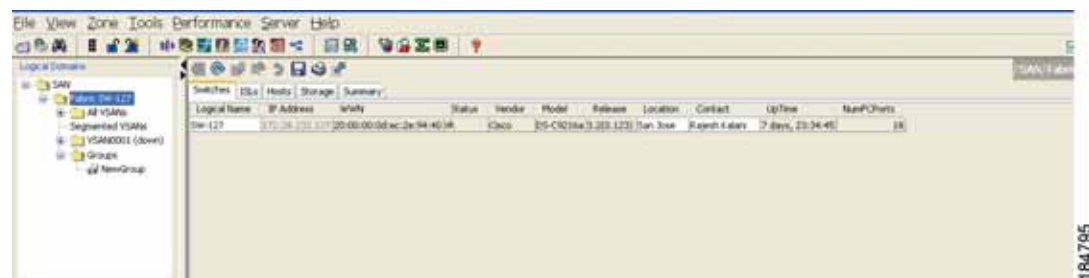
Step 3 Click **OK** to save the updated group.

To filter the display by a group you have created, perform this task:

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

You see the list of groups that you have created as shown in [Figure 5-10](#).

Figure 5-10 Group Highlighted in Fabric Pane Map



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

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

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



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

Send comments to nx5000-docfeedback@cisco.com

Status Bar

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

Setting Fabric Manager Preferences

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

- General
- SNMP
- Map

The default General preferences for Fabric Manager are as follows:

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



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

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

Send comments to nx5000-docfeedback@cisco.com

The default SNMP preferences for Fabric Manager are as follows:

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

The default Map preferences for Fabric Manager are as follows:

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

Network Fabric Discovery

Cisco Fabric Manager collects information about the fabric topology through SNMP queries to the switches that are connected to Fabric Manager. The switch replies after having discovered all devices connected to the fabric by using the information from its FSPF technology database and the Name Server

Send comments to nx5000-docfeedback@cisco.com

database and collected using the Fabric Configuration Server's request/response mechanisms that are defined by the FC-GS-3/4 standard. When you start Fabric Manager, you enter the IP address (or host name) of a seed switch for discovery.

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

Modifying the Device Grouping

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

To group end devices in a single enclosure to have them represented by a single icon on the map, Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **End Devices** and then choose **Storage** or **Hosts**.
You see the end devices displayed in the Information pane.
 - Step 2** Click one of the devices in the Fabric pane, or click the **Enclosures** tab of the Information pane, and then click the device name (in the Name field) that you want to include in the enclosure.
 - Step 3** Enter a name to identify the new enclosure in the Fabric pane map.
 - Step 4** Click once on the device name in the Name field. To choose more than one name, press the **Shift** key and click each of the other names.
 - Step 5** Press **Ctrl-C** to copy the selected name(s).
 - Step 6** Press **Ctrl-V** to paste the device name into the Name field.



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

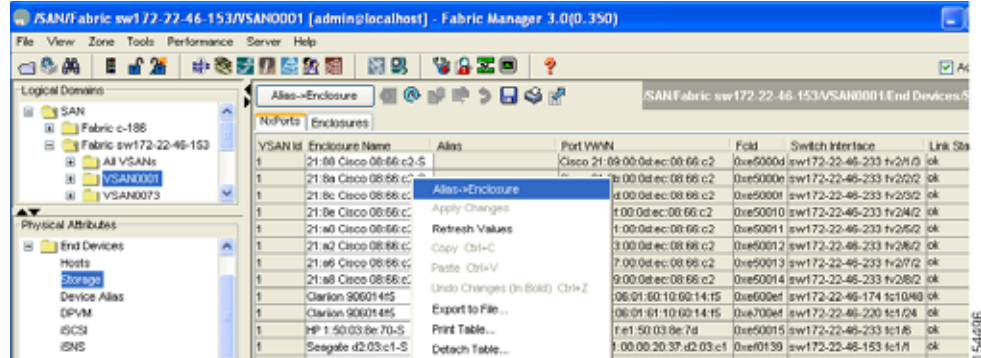
Using Alias Names as Enclosures

To create an enclosure that uses the alias name as the name of the enclosure using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **End Devices** and choose **Hosts** or **Storage**.
You see the list of devices in the Information pane. The NxPorts tab is the default.
 - Step 2** Right-click the enclosure names that you want to convert to alias names and choose **Alias > Enclosure** as shown in [Figure 5-11](#).

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Figure 5-11 Alias Enclosure



Step 3 Click the **Apply Changes** icon to save these changes.

Controlling Administrator Access with Users and Roles

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

The default-role contains the access permissions needed by a user to access the GUI (Fabric Manager and Device Manager). These access permissions are automatically granted to all users in order for them to use the GUI.

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

Using Fabric Manager Wizards

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

- **VSAN**—Creates VSANs on multiple switches in the fabric and sets VSAN attributes including interop mode, load balancing, and FICON.
- **Zone Edit Tool**—Creates zone sets, zones, and aliases. Adds members to zones and edits the zone database.
- **SAN Port Channel**—Creates SAN port channels from selected ISLs either manually or automatically. Sets SAN port channel attributes such as channel ID and trunking mode.
- **IP ACL**—Creates ordered IP access control lists and distributes to selected switches in the fabric.
- **License Install**—Facilitates download and installation of licenses in selected switches in the fabric.
- **Software Install**—Verifies image compatibility and installs software images on selected switches in the fabric.
- **Virtual Interface Group**—Creates virtual interface groups (VIGs).

Send comments to nx5000-docfeedback@cisco.com

- Virtual Interfaces—Creates virtual Ethernet interfaces and virtual Fibre Channel interfaces.

Fabric Manager Troubleshooting Tools

Fabric Manager has several troubleshooting tools available from the toolbar or Tools menu. Procedures for using these tools are described in [Chapter 31, “Troubleshooting Your Fabric.”](#) This section provides a brief description of each tool:

- Zone Merge Analysis—The zone merge analysis tool (available from the Zone menu) allows you to determine if zones will merge successfully when two Cisco switches are interconnected. If the interconnected switch ports allow VSANs with identical names or contain zones with identical names, then Fabric Manager verifies that the zones contain identical members. The merge analysis tool can be run before attempting a merge or after fabrics are interconnected to determine zone merge failure causes.
- End-to-End Connectivity—Fabric Manager’s end-to-end connectivity analysis tool uses FC Ping to verify interconnections between Cisco switches and end-device (HBAs and storage devices) in a particular VSAN. End devices must be connected to a manageable switch through an active in-band or out-of-band management path. In addition to basic connectivity, Fabric Manager can optionally verify the following:
 - Paths are redundant.
 - Zones contain at least two members.

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

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

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



CHAPTER 6

Device Manager

This chapter describes Cisco Device Manager and provides procedures for using it. This chapter contains the following sections:

- [Information About Device Manager, page 6-1](#)
- [Launching Device Manager, page 6-2](#)
- [Using Device Manager, page 6-2](#)
- [Using the Quick Configuration Tool, page 6-6](#)
- [Setting Device Manager Preferences, page 6-7](#)

Information About Device Manager

Device Manager provides a graphic representation of a Cisco Cisco Nexus 5000 Series switch chassis, including the installed expansion modules, the status of each port, the power supplies, and the fan assemblies.

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

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

- Configure zones for multiple VSANs.
- Manage ports, SAN port channels, and trunking.
- Manage SNMPv3 security access to switches.
- Manage CLI security access to the switch.
- Manage alarms, events, and notifications.
- Save and copy configuration files and software image.
- View hardware configuration.
- View chassis, module, port status, and statistics.

Send comments to nx5000-docfeedback@cisco.com

Launching Device Manager

You can launch Device Manager either from your desktop or from Fabric Manager.

To launch Device Manager from your desktop, double-click the **Device Manager** icon and follow the instructions described in the “[Integrating Cisco Fabric Manager with Other Management Tools](#)” section on page 2-18.

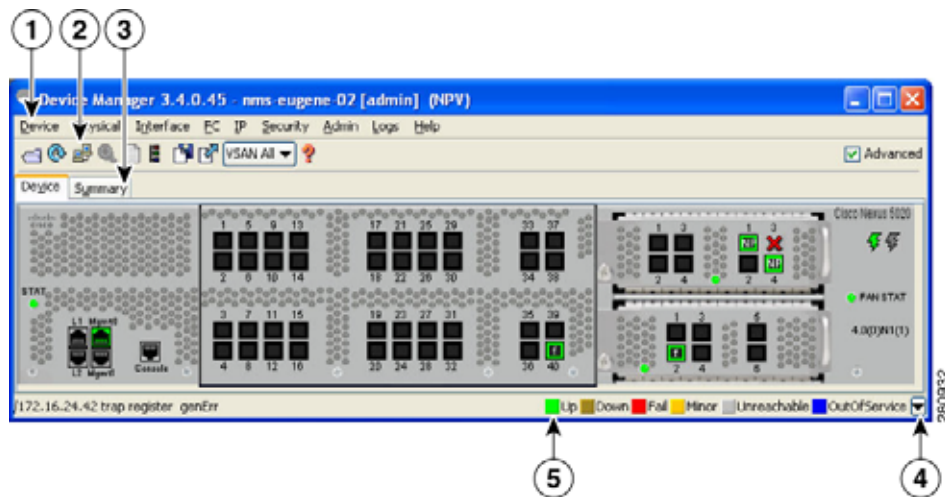
To launch Device Manager from Fabric Manager, perform one of these tasks:

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

Using Device Manager

This section describes the Device Manager interface as shown in [Figure 6-1](#).

Figure 6-1 Device Manager, Device Tab



1	Menu bar	4	Legend
2	Toolbar	5	Status
3	Tabs		

Menu Bar

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

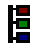




Send comments to nx5000-docfeedback@cisco.com

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

Toolbar Icons







The Device Manager toolbar provides quick access to many Device Manager features. Once the icon is selected, a dialog box may open that allows configuration of the feature. The toolbar provides the main Device and Summary View icons as shown in [Table 6-1](#).

Table 6-1 Device Manager Main Toolbar

Icon	Description
 Open Device	Opens the Device Manager view for another switch, with the option to open this view in a separate window.
 Refresh Display	Communicates with the switch and displays the information in the Device Manager view.
 Command-Line Interface	Opens a separate CLI command window to the switch.
 Configure Selected	Opens a configuration dialog box for the selected component (line card or port).
 SysLog	Opens a window that lists the latest system messages that occurred on the switch.

Send comments to nx5000-docfeedback@cisco.com

Table 6-1 Device Manager Main Toolbar (continued)

Icon	Description
 VSANs	Opens the VSAN dialog box that provides VSAN configuration for the switch.
 Save Configuration	Saves the current running configuration to the startup configuration.
 Copy	Copies configuration file between server and switch
 Toggle FICON/Interface Port Labels	Toggles the FICON and interface port labels.
 Select VSAN	Filters the port display to show only those ports belonging to the selected VSAN.
 Help	Accesses online help for Device Manager.

Dialog Boxes

If a toolbar icon is selected, a dialog box may open that allows configuration of the selected feature. The dialog box may include table manipulation icons. See the “[Information Pane](#)” section on page 5-11 for descriptions of these icons.

Tabs

Tabs provide the following functions:

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

Legend

The legend at the bottom right of the Device Manager indicates the following port status:

- Colors
 - Green—The port is up.

Send comments to nx5000-docfeedback@cisco.com

- Brown—The port is administratively down.
- Red—The port is down or has failed.
- Amber—The port has a minor fault condition.
- Gray—The port is unreachable.
- Blue—The port is out of service.
- Labels
 - X—Link failure
 - E—ISL
 - TE—Multi-VSAN ISL
 - F—Host/storage
 - FL—F loop
 - I— iSCSI
 - SD—SPAN destination
 - CH—Channel
 - CU—Control unit
 - NP - Proxy N port (NPV mode)
 - f—virtual Fibre Channel interface is present
 - e—virtual Ethernet interface is present

Supervisor and Switching Modules

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



Tip

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

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



Note

In Device Manager, Enable and Disable are available only for Fibre Channel ports (and not Ethernet ports).

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

Send comments to nx5000-docfeedback@cisco.com

To create SAN port channels using Device Manager, click **PortChannels** from the Interface menu. For detailed instructions, see [Chapter 14, “Configuring SAN Port Channels.”](#) You can also use Fabric Manager to conveniently create a SAN port channel.



Note

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

Context Menus

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

Device View menus:

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



Note

Context menus for Ethernet ports configure and monitor the physical Ethernet port in addition to any virtual interfaces that exist on the physical Ethernet.

Summary View menus:

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

Using the Quick Configuration Tool

Device Manager provides a tool for configuring the Ethernet interfaces on the switch. The Quick Configuration Tool allows you to select one of the following configurations for each Ethernet interface:

- **Eth Only**—Configures the physical Ethernet without any virtual interfaces.
- **vEth Only**—Configures the physical Ethernet to have an associated VIG with a virtual Ethernet interface.
- **vFC Only**—Configures the physical Ethernet to have an associated VIG with a virtual Fibre Channel interface.
- **vEth + vFC**—Configures the physical Ethernet to have an associated VIG with a virtual Ethernet interface and a virtual Fibre Channel interface.

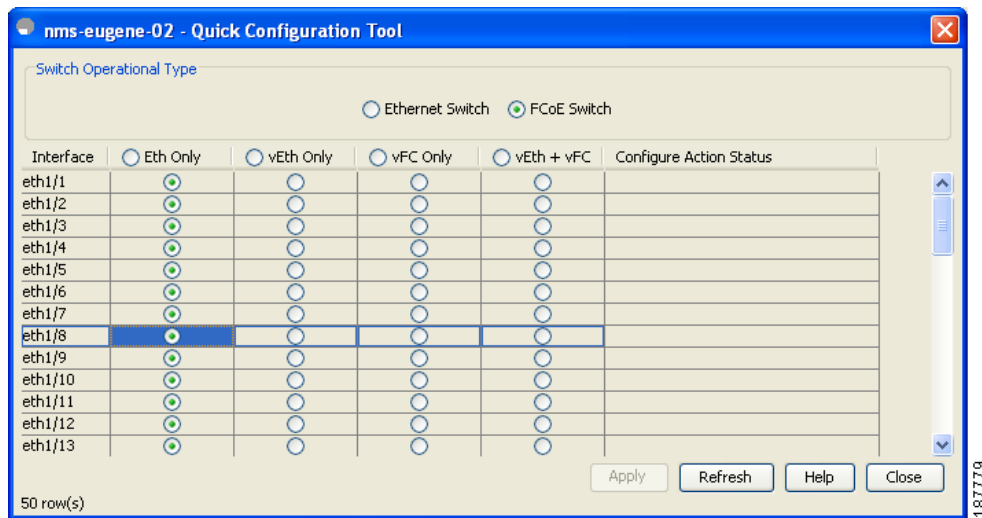
To configure the Ethernet interfaces using the Quick Configuration Tool, perform this task:

Step 1 In the tools menu, choose Quick Config Wizard.

You see the Quick Configuration Tool window as shown in [Figure 6-2](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 6-2 Quick Configuration Tool



- Step 2** In the Switch Operational Type pane, click the **Ethernet Switch** radio button if you are configuring the switch as a pure Ethernet switch. Click the **FCoE Switch** radio button if you are configuring the switch as an I/O consolidation switch (with Fibre Channel and FCoE interfaces).
- Step 3** (Optional) Click the button in the column header to set all of the interfaces to the value in the selected column.
- Step 4** For each row, click the radio button for the configuration you want to apply to this interface.
- Step 5** Click the **Apply** button to apply the configuration changes.
- The Configure Action Status field displays the current status of the requested configuration change.
- Step 6** (Optional) Click the **Refresh** button to clear the Configure Action Status field. Device Manager then updates the field with the latest status from the switch.

Setting Device Manager Preferences

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

- **Retry Requests x Time(s) After x sec Timeout**—Allows you to set the retry request values. The default settings are 1 time after a 5-second timeout.
- **Enable Status Polling Every x secs**—Allows you to set the status polling value. The default setting is enabled (checked) with a time of 40 seconds.
- **Trace SNMP Packets in Message Log**—Allows you to set whether Device Manager traces SNMP packets and logs the trace. The default setting is disabled (unchecked).
- **Register for Events After Open, Listen on Port 1163**—Allows you to register this switch so that events are logged once you open Device Manager. The default setting is enabled (checked).
- **Confirm Deletion**—Displays a pop-up confirmation when you delete part of your configuration using Device Manager. The default setting is enabled (checked).

Send comments to nx5000-docfeedback@cisco.com

- **Show WorldWideName (WWN) Vendor**—Displays the world wide name vendor name in any table or listing displayed by Device Manager. If **Prepend** is checked, the name is displayed in front of the IP address of the switch. If **Replace** is checked, the name is displayed instead of the IP address. The default setting is enabled (checked) with the **Prepend** option.
- **Show Timestamps as Date/Time**—Displays timestamps in the date/time format. If this preference is not checked, timestamps are displayed as elapsed time. The default setting is enabled (checked).
- **Telnet Path**—Sets the path for the telnet.exe file on your system. The default is telnet.exe, but you need to browse for the correct location.



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

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



CHAPTER 7

Using Cisco Fabric Services

Cisco Nexus 5000 Series switches provide Cisco Fabric Services (CFS) capability, which simplifies provisioning by automatically distributing configuration information to all switches in the network.

Switch features can use the CFS infrastructure to distribute feature data or configuration data required by the feature.

This chapter contains the following sections:

- [Information About CFS, page 7-1](#)
- [CFS Distribution, page 7-2](#)
- [CFS Support for Applications, page 7-6](#)
- [CFS Regions, page 7-10](#)
- [Displaying CFS Distribution Information, page 7-15](#)
- [CFS Example Using Fabric Manager, page 7-15](#)
- [CFS Example Using Device Manager, page 7-18](#)
- [Default Settings, page 7-19](#)

Information About CFS

Some features in the Cisco Nexus 5000 Series switch require configuration synchronization with other switches in the network to function correctly. Synchronization through manual configuration at each switch in the network can be a tedious and error-prone process.

Cisco Fabric Services (CFS) provides a common infrastructure for automatic configuration synchronization in the network. It provides the transport function and a set of common services to the features. CFS has the ability to discover CFS capable switches in the network and discovering feature capabilities in all CFS capable switches.

Cisco Nexus 5000 Series switches support CFS message distribution over Fibre Channel, IPv4 or IPv6 networks. If the switch is provisioned with Fibre Channel ports, CFS over Fibre Channel is enabled by default. CFS over IP must be explicitly enabled.

CFS provides the following features:

- Peer-to-peer protocol with no client-server relationship at the CFS layer.
- CFS message distribution over Fibre Channel, IPv4 or IPv6 networks.
- Three modes of distribution.

Send comments to nx5000-docfeedback@cisco.com

- Coordinated distributions: Only one distribution is allowed in the network at any given time.
- Uncoordinated distributions: Multiple parallel distributions are allowed in the network except when a coordinated distribution is in progress.
- Unrestricted uncoordinated distributions: Multiple parallel distributions are allowed in the network in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.

The following features are supported for CFS distribution over IP:

- One scope of distribution over an IP network:
 - Physical scope: The distribution spans the entire IP network.

The following features are supported for CFS distribution over Fibre Channel SANs:

- Three scopes of distribution over SAN fabrics.
 - Logical scope: The distribution occurs within the scope of a VSAN.
 - Physical scope: The distribution spans the entire physical topology.
 - Over a selected set of VSANs: Some features require configuration distribution over some specific VSANs. These features can specify to CFS the set of VSANs over which to restrict the distribution.
- Supports a merge protocol that facilitates the merge of feature configuration during a fabric merge event (when two independent SAN fabrics merge).

CFS Distribution

The CFS distribution functionality is independent of the lower layer transport. Cisco Nexus 5000 Series switches support CFS distribution over IP and CFS distribution over Fibre Channel. Features that use CFS are unaware of the lower layer transport.

Additional details are provided in the following sections:

- [CFS Distribution Modes, page 7-2](#)
- [Enabling/Disabling CFS Distribution on a Switch, page 7-3](#)
- [CFS Distribution over IP, page 7-4](#)
- [CFS Distribution over Fibre Channel, page 7-5](#)
- [CFS Distribution Scopes, page 7-5](#)
- [CFS Merge Support, page 7-6](#)

CFS Distribution Modes

CFS supports three distribution modes to accommodate different feature requirements. Only one mode is allowed at any given time. CFS distribution modes are described in the following sections:

- [Uncoordinated Distribution, page 7-3](#)
- [Coordinated Distribution, page 7-3](#)
- [Unrestricted Uncoordinated Distributions, page 7-3](#)

Send comments to nx5000-docfeedback@cisco.com

Uncoordinated Distribution

Uncoordinated distributions are used to distribute information that is not expected to conflict with that from a peer. Parallel uncoordinated distributions are allowed for a feature.

Coordinated Distribution

Coordinated distributions allow only one feature distribution at a given time. CFS uses locks to enforce this. A coordinated distribution is not allowed to start if locks are taken for the feature anywhere in the network. A coordinated distribution consists of three stages:

1. A network lock is acquired.
2. The configuration is distributed and committed.
3. The network lock is released.

Coordinated distribution has two variants:

- CFS driven—The stages are executed by CFS in response to a feature request without intervention from the feature.
- Feature driven—The stages are under the complete control of the feature.

Coordinated distributions are used to distribute information that can be manipulated and distributed from multiple switches, for example, the port security configuration.

Unrestricted Uncoordinated Distributions

Unrestricted uncoordinated distributions allow multiple parallel distributions in the network in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.

Enabling/Disabling CFS Distribution on a Switch

If the switch is provisioned with Fibre Channel ports, CFS over Fibre Channel is enabled by default. CFS over IP must be explicitly enabled.

You can globally disable CFS on a switch to isolate the features using CFS from network-wide distributions while maintaining physical connectivity. When CFS is globally disabled on a switch, CFS operations are restricted to the switch.

To globally disable or enable CFS distribution on a switch using Fabric Manager, perform this task:

-
- Step 1** Choose any CFS feature. For example, expand **Switches > Events** and then choose **CallHome** in the Physical Attributes pane.
The Information pane shows that feature, with a CFS tab.
 - Step 2** Click the **CFS** tab to display the CFS state for each switch in the network for that feature.
 - Step 3** Click a value in the **Global State** column. The value changes to a drop-down menu.
 - Step 4** From the drop-down menu, choose **disable** or **enable**.
 - Step 5** Repeat steps 3 and 4 for all switches that you want to disable or enable CFS.
 - Step 6** Set the Config Action column to **commit**.

Send comments to nx5000-docfeedback@cisco.com

- Step 7** Click the **Apply Changes** icon to commit the configuration changes for that feature and distribute the changes through CFS.

To globally disable or enable CFS distribution on a switch using Device Manager, perform this task:

- Step 1** Choose **Admin > CFS (Cisco Fabric Services)**.
You see the CFS dialog box with the CFS status for all features on that switch.
- Step 2** Uncheck or check the **Globally Enabled** check box to disable or enable CFS distribution on this switch.
- Step 3** Click **Apply** to disable CFS on this switch.

CFS Distribution over IP

CFS distribution over IP supports the following features:

- Physical distribution over an entirely IP network.
- Physical distribution over a hybrid Fibre Channel and IP network with the distribution reaching all switches that are reachable over either Fibre Channel or IP.



Note The switch attempts to distribute information over Fibre Channel first and then over the IP network if the first attempt over Fibre Channel fails. CFS does not send duplicate messages if distribution over both IP and Fibre Channel is enabled.

- Distribution over IP version 4 (IPv4) or IP version 6 (IPv6).




Note CFS cannot distribute over both IPv4 and IPv6 from the same switch.

- Keepalive mechanism to detect network topology changes using a configurable multicast address.
- Compatibility with Cisco MDS 9000 Family switches running release 2.x or later.

Figure 7-1 shows a network with both Fibre Channel and IP connections. Node A forwards an event to node B over Fibre Channel. Node B forwards the event node C and node D using unicast IP. Node C forwards the event to node E using Fibre Channel.

Figure 7-1 Network Example 1 with Fibre Channel and IP Connections



[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Figure 7-2 is the same as Figure 7-1 except that node C and node D are connected using Fibre Channel. All processes is the same in this example because node B has node C and node D the distribution list for IP. Node C does not forward to node D because node D is already in the distribution list from node B.

Figure 7-2 Network Example 2 with Fibre Channel and IP Connections





Figure 7-3 is the same as Figure 7-2 except that node D and node E are connected using IP. Both node C and node D forward the event to E because the node E is not in the distribution list from node B.

Figure 7-3 Network Example 3 with Fibre Channel and IP Connections



CFS Distribution over Fibre Channel

For FCS distribution over Fibre Channel, the CFS protocol layer resides on top of the FC2 layer. CFS uses the FC2 transport services to send information to other switches. CFS uses a proprietary SW_ILS (0x77434653) protocol for all CFS packets. CFS packets are sent to or from the switch domain controller addresses.

CFS Distribution Scopes

Different applications on the Cisco Nexus 5000 Series switches need to distribute the configuration at various levels. The following levels are available when using CFS distribution over Fibre Channel:

- VSAN level (logical scope)

Applications that operate within the scope of a VSAN have the configuration distribution restricted to the VSAN. An example application is port security where the configuration database is applicable only within a VSAN.



Note Logical scope is not supported for FCS distribution over IP.

Send comments to nx5000-docfeedback@cisco.com

- Physical topology level (physical scope)
Some applications (such as NTP) need to distribute the configuration to the entire physical topology.
- Between two selected switches
Some applications operate only between selected switches in the network.

CFS Merge Support

CFS Merge is supported for CFS distribution over Fibre Channel.

An application keeps the configuration synchronized in a SAN fabric through CFS. Two such fabrics might merge as a result of an ISL coming up between them. These two fabrics could have two different sets of configuration information that need to be reconciled in the event of a merge. CFS provides notification each time an application peer comes online. If a fabric with M application peers merges with another fabric with N application peers, and if an application triggers a merge action on every notification, a link-up event results in M*N merges in the fabric.

CFS supports a protocol that reduces the number of merges required to one by handling the complexity of the merge at the CFS layer. This protocol runs per application per scope. The protocol involves selecting one switch in a fabric as the merge manager for that fabric. The other switches do not have a role in the merge process.

During a merge, the merge manager in the two fabrics exchange their configuration databases with each other. The application on one of them merges the information, decides if the merge is successful, and informs all switches in the combined fabric of the status of the merge.

In case of a successful merge, the merged database is distributed to all switches in the combined fabric and the entire new fabric remains in a consistent state. You can recover from a merge failure by starting a distribution from any of the switches in the new fabric. This distribution restores all peers in the fabric to the same configuration database.

CFS Support for Applications

The following topics describe the CFS capabilities that support applications:

- [CFS Application Requirements, page 7-6](#)
- [Enabling CFS for an Application, page 7-7](#)
- [Locking the Network, page 7-8](#)
- [Committing Changes, page 7-8](#)
- [Discarding Changes, page 7-9](#)
- [Saving the Configuration, page 7-10](#)
- [Clearing a Locked Session, page 7-10](#)

CFS Application Requirements

All switches in the network must be CFS capable. Switches that are not CFS capable do not receive distributions and result in part of the network not receiving the intended distribution.

Send comments to nx5000-docfeedback@cisco.com

CFS has the following requirements:

- **Implicit CFS usage**—The first time you issue a CFS task for a CFS-enabled application, the configuration modification process begins and the application locks the network.
- **Pending database**—The pending database is a temporary buffer to hold uncommitted information. The uncommitted changes are not applied immediately to ensure that the database is synchronized with the database in the other switches in the network. When you commit the changes, the pending database overwrites the configuration database (also known as the active database or the effective database).
- **CFS distribution enabled or disabled on a per-application basis**—The default (enable or disable) for CFS distribution state differs between applications. If CFS distribution is disabled for an application, then that application does not distribute any configuration nor does it accept a distribution from other switches in the network.
- **Explicit CFS commit**—Most applications require an explicit commit operation to copy the changes in the temporary buffer to the application database, to distribute the new database to the network, and to release the network lock. The changes in the temporary buffer are not applied if you do not perform the commit operation.


Enabling CFS for an Application

All CFS-based applications provide an option to enable or disable the distribution capabilities.

Applications have the distribution enabled by default.

The application configuration is not distributed by CFS unless distribution is explicitly enabled for that application.

To enable CFS for a feature using Fabric Manager, perform this task:

-
- Step 1** Choose a feature on which to enable CFS.
- For example, expand **Switches > Events** and then choose **CallHome** in the Physical Attributes pane. The Information pane shows that feature with a CFS tab. Click the **CFS** tab to display the CFS state for each switch in the network for that feature.
- Step 2** Decide on which switches to enable CFS. Set the Feature Admin column to either **enable** to enable CFS or **disable** to disable CFS.
-  **Note** Enable CFS for all switches in the network or VSAN for the feature that uses CFS.
-
- Step 3** Right-click the row you changed to see the pop-up menu. Choose **Apply Changes** to apply the CFS configuration change. The CFS tab updates as the CFS changes take effect.
- Fabric Manager retrieves the status of the CFS change and updates the Last Result column.
-

Send comments to nx5000-docfeedback@cisco.com

To enable CFS for a feature using Device Manager, perform this task:

Step 1 Choose **Admin > CFS (Cisco Fabric Services)**.

You see the CFS dialog box with the CFS status for all features on that switch.

Step 2 Decide which features need CFS. Set the Command column to either **enable** to enable CFS or **disable** to disable CFS.



Note Enable or disable CFS for all switches in the network or VSAN for the feature that uses CFS.

Step 3 Click **Pending Differences** to compare the configuration of this feature on this switch to other switches in the network or VSAN that have CFS enabled for this feature. Close the Show Pending Diff dialog box.

Step 4 Click **Apply** to apply the CFS configuration change.

Device Manager retrieves the status of the CFS change and updates the Last Command and Result columns.

Locking the Network

When you configure (first time configuration) a feature (or application) that uses the CFS infrastructure, that feature starts a CFS session and locks the network. When a network is locked, the switch software allows configuration changes to this feature only from the switch holding the lock. If you make configuration changes to the feature from another switch, the switch issues a message to inform the user about the locked status. The configuration changes are held in a pending database by that application.

If you start a CFS session that requires a network lock but forget to end the session, an administrator can clear the session. If you lock a network at any time, your user name is remembered across restarts and switchovers. If another user (on the same machine) tries to perform configuration tasks, that user's attempts are rejected.

Committing Changes

A commit operation saves the pending database for all application peers and releases the lock for all switches.

In general, the commit function does not start a session, only a lock function starts a session. However, an empty commit is allowed if configuration changes are not previously made. In this case, a commit operation results in a session that acquires locks and distributes the current database.

When you commit configuration changes to a feature using the CFS infrastructure, you receive a notification about one of the following responses:

- One or more external switches report a successful status—The application applies the changes locally and releases the network lock.
- None of the external switches report a successful state—The application considers this state a failure and does not apply the changes to any switch in the network. The network lock is not released.

You can commit changes for a specified feature by for that feature.

Send comments to nx5000-docfeedback@cisco.com

To commit changes using Fabric Manager for CFS-enabled features, perform this task:

-
- Step 1** Choose the feature you want to enable CFS for.
- For example, expand **Switches** expand **Events**, and then choose **CallHome** from the Physical Attributes pane.
- The Information pane shows that feature with a CFS tab.
- Step 2** Click the **CFS** tab to display the CFS state for each switch in the network for that feature.
- Step 3** Right-click the value in the Config Action column for any switch and choose an option from the drop-down menu (Copy, Paste, Export to File, Print Table, Detach Table).
- Step 4** Click the **Apply Changes** icon to commit the configuration changes for that feature and distribute the changes through CFS.
- Fabric Manager retrieves the status of the CFS change and updates the Last Command and Last Result columns for the feature or VSAN.
-

To commit changes using Device Manager for CFS-enabled features, perform this task:

-
- Step 1** Choose **Admin > CFS (Cisco Fabric Services)**.
- You see the CFS dialog box with the CFS status for all features on that switch.
- Step 2** For each applicable feature, set the Command column to **commit** to commit the configuration changes for that feature and distribute the changes through CFS, or set it to **abort** to discard the changes for that feature and release the network lock for CFS for that feature.
- Step 3** Optionally, provide a **Type** or **VsanID** as the basis for the CFS distribution for CFS features that require this information.
- Step 4** Click **Pending Differences** to check the configuration of this feature on this switch as compared to other switches in the network or VSAN that have CFS enabled for this feature.
- Step 5** Click **Apply** to apply the CFS configuration change.
- Device Manager retrieves the status of the CFS change and updates the Last Command and Result columns.
-



Caution

If you do not commit the changes, they are not saved to the running configuration.

Discarding Changes

If you discard configuration changes, the application flushes the pending database and releases locks in the network. Both the abort and commit functions are only supported from the switch from which the network lock is acquired.

You can discard changes for a specified feature by setting the Command column value to **disable** for that feature then clicking **Apply**.

Send comments to nx5000-docfeedback@cisco.com

Saving the Configuration

Configuration changes that have not been applied yet (still in the pending database) are not shown in the running configuration. The configuration changes in the pending database overwrite the configuration in the effective database when you commit the changes.



Caution

If you do not commit the changes, they are not saved to the running configuration.

The CISCO-CFS-MIB contains SNMP configuration information for any CFS-related functions. Refer to the *Cisco Cisco Nexus 5000 Series MIB Quick Reference* for more information on this MIB.

Clearing a Locked Session

You can clear locks held by an application from any switch in the network to recover from situations where locks are acquired and not released. This function requires Admin permissions.

To clear locks using Fabric Manager, perform this task:

- Step 1 Click the **CFS** tab.
- Step 2 Choose **clearLock** from the Config Action drop-down list for each switch that you want to clear the lock (see [Figure 7-4](#)).
- Step 3 Click the **Apply Changes** icon to save the change.

Figure 7-4 Clearing Locks

Switch	Feature Admin	Feature Oper	Global State	Config Action	Last Command	Last Result	Lock Owner Switch	Lock Owner User Name	Merge Status	Master	Scope
sw172-22-46-221	noSelection	enabled	enable	noSelection					success	<input type="checkbox"/>	fFabric: p1network
sw172-22-46-220	noSelection	enabled	enable	noSelection	commitChanges	success	sw172-22-46-220	neprivate	success	<input checked="" type="checkbox"/>	fFabric: p1network
sw172-22-46-174	noSelection	enabled	enable	noSelection					success	<input type="checkbox"/>	fFabric: p1network



Caution

Exercise caution when using this function to clear locks in the network. Any pending configurations in any switch in the network is flushed and lost.

CFS Regions

This section contains the following topics:

- [About CFS Regions, page 7-11](#)
- [Example Scenario, page 7-11](#)
- [Managing CFS Regions Using Fabric Manager, page 7-11](#)
- [Creating CFS Regions, page 7-12](#)
- [Assigning Features to CFS Regions, page 7-12](#)

Send comments to nx5000-docfeedback@cisco.com

- [Moving a Feature to a Different Region, page 7-13](#)
- [Removing a Feature from a Region, page 7-14](#)
- [Deleting CFS Regions, page 7-14](#)

About CFS Regions

A CFS region is a user-defined subset of switches for a given feature or application in its physical distribution scope. When a network spans a vast geography, you may need to localize or restrict the distribution of certain profiles among a set of switches based on their physical proximity. CFS regions allow you to create multiple islands of distribution within the network for a given CFS feature or application. CFS regions are designed to restrict the distribution of a feature's configuration to a specific set or grouping of switches in a network.



Note You can only configure a CFS region based on physical switches. You cannot configure a CFS region in a VSAN.

Example Scenario

The Call Home application triggers alerts to network administrators when a situation arises or something abnormal occurs. When the network covers many geographies, and there are multiple network administrators who are each responsible for a subset of switches in the network, the Call Home application sends alerts to all network administrators regardless of their location. For the Call Home application to send message alerts selectively to network administrators, the physical scope of the application has to be fine tuned or narrowed down, which is achieved by implementing CFS regions.

CFS regions are identified by numbers ranging from 0 through 200. Region 0 is reserved as the default region, and contains every switch in the network. You can configure regions from 1 through 200. The default region maintains backward compatibility.

If the feature is moved, that is, assigned to a new region, its scope is restricted to that region; it ignores all other regions for distribution or merging purposes. The assignment of the region to a feature has precedence in distribution over its initial physical scope.

You can configure a CFS region to distribute configurations for multiple features. However, on a given switch, you can configure only one CFS region at a time to distribute the configuration for a given feature. Once you assign a feature to a CFS region, its configuration cannot be distributed within another CFS region.

Managing CFS Regions Using Fabric Manager

This section describes how to use Fabric Manager for managing CFS regions. Fabric Manager provides a comprehensive view of all the switches, regions, and the features associated with each region in the topology. To complete the following tasks, use the tables under the All Regions and Feature by Region tabs:

- [Creating CFS Regions, page 7-12](#)
- [Assigning Features to CFS Regions, page 7-12](#)
- [Moving a Feature to a Different Region, page 7-13](#)

Send comments to nx5000-docfeedback@cisco.com


- [Removing a Feature from a Region, page 7-14](#)

Creating CFS Regions

To create a CFS region using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches** and then choose **CFS**.
The information pane displays the Global, IP Multicast, Feature by Region, and All Regions tabs.
- Step 2** Click the **All Regions** tab.
The tab displays a list of Switches and RegionIds.
- Step 3** Click the **Create Row** button on the toolbar.
[Figure 7-5](#) shows the Create a Region dialog box.

Figure 7-5 Create a Region Dialog Box



- Step 4** Choose the switch from the drop-down list and choose a RegionId from the range.
- Step 5** Click **Create**.
Upon successful creation of the region, Success is displayed at the bottom of the dialog box.
-


Assigning Features to CFS Regions

To assign a feature to a region using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches** and then choose **CFS**.
The information pane displays the Global, IP Multicast, Feature by Region, and All Regions tabs.
- Step 2** Click the **Feature by Region** tab.
This tab lists all the switches along with their corresponding Feature and RegionId.
- Step 3** Click the **Create Row** button on the toolbar.
[Figure 7-6](#) shows the Assign a Feature dialog box.

Send comments to nx5000-docfeedback@cisco.com

Figure 7-6 Assign a Feature Dialog Box



- Step 4** Choose a switch from the drop-down box.
The features running on the selected switch are listed in the Feature drop-down box.
- Step 5** Choose a feature on that switch to associate a region.
- Step 6** Choose the region number from the list to associate a RegionId with the selected feature.
- Step 7** Click **Create** to complete assignment of a switch feature to the region.
Upon successful assignment of feature, Success is displayed at the bottom of the dialog box.

When a feature is assigned to a new region using the Feature by Region tab, a new row with the new region is created automatically in the table under the All Regions tab. Alternatively, you can create a region using the All Regions tab.



Note

In the Feature by Region tab, when you try to reassign a feature on a switch to another region by clicking **Create Row**, an operation failed message is shown. The error message states that an entry already exists. However, moving a feature to a different region is a different task and it is described in the next section.

Moving a Feature to a Different Region


Before moving a feature to a new region, create the new region in the All Regions tab. That is, a new row has to be added in the All Regions tab with the new Region ID.

To move a feature to a different region using Fabric Manager, perform this task:

- Step 1** In the Physical Attributes pane, expand **Switches** and then choose **CFS**.
The information pane displays the Global, IP Multicast, Feature by Region, and All Regions tabs.
- Step 2** Click the **Feature by Region** tab.
[Figure 7-7](#) shows the Feature by Region tab, which lists all the switches along with their feature and region details.

Send comments to nx5000-docfeedback@cisco.com

Figure 7-7 Feature by Region Tab



- Step 3 Double-click the RegionId cell in the required row.
The cursor blinks in the cell prompting a change in the value.
- Step 4 Change the RegionId value to the required region.
- Step 5 Click the **Apply Changes** button on the tool bar to commit the change.

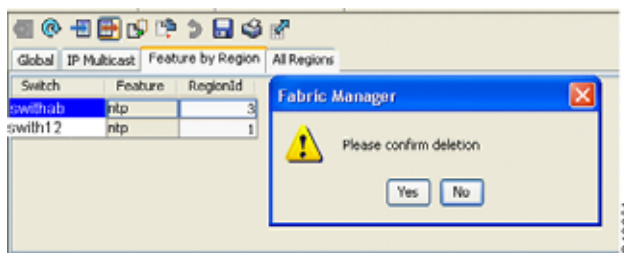
Removing a Feature from a Region

To remove a feature from a region using Fabric Manager, perform this task:

- Step 1 Click the **Feature by Region** tab and click the required row.
- Step 2 Click the **Delete Row** button on the toolbar.

Figure 7-8 shows a confirmation dialog box.

Figure 7-8 Removing a Feature from a Region



- Step 3 Click **Yes** to confirm row deletion from the table in view.

Deleting CFS Regions

To delete an entire region, perform this task:

- Step 1 Click the **All Regions** tab and click the required row.
- Step 2 Click **Delete Row**.

This action removes all entries pertaining to that switch and region in the table under Feature by Region tab.

Send comments to nx5000-docfeedback@cisco.com

Figure 7-9 shows a confirmation dialog box.

Figure 7-9 Deleting CFS Regions



Step 3 Click **Yes** to confirm deletion of the region.



Note CFS cannot distribute over both IPv4 and IPv6 from the same switch.

Displaying CFS Distribution Information

To display the status of CFS distribution on the switch using Device Manager, perform this task:

Step 1 Choose **Admin > CFS (Cisco Fabric Services)**.

You see the CFS dialog box. This dialog box displays the distribution status of each feature using CFS, which currently registered applications are using CFS, and the result of the last successful merge attempt.

Step 2 Select a row and click **Details** to view more information about the feature.

CFS Example Using Fabric Manager

This procedure is an example of what you see when you use Fabric Manager to configure a feature that uses CFS:

Step 1 Select the CFS-capable feature that you want to configure.

For example, expand a **VSAN** and then choose **Port Security** in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **CFS** tab.

You see the CFS configuration and status for each switch (see [Figure 7-10](#)).

Send comments to nx5000-docfeedback@cisco.com

Figure 7-10 CFS Configuration

Switch	Feature Admin	Feature Oper	Global State	Config Action	Last Command	Last Result	Lock Owner Switch	Lock Owner User Name	Merge Status	Master	Scope
sw172-22-46-220	noSelection	enabled	enable	noSelection			sw172-22-46-220	new	success	<input checked="" type="checkbox"/>	vsanScope
sw172-22-46-174	noSelection	enabled	enable	noSelection					success	<input type="checkbox"/>	vsanScope
sw172-22-46-221	noSelection	disabled	enable	noSelection						<input type="checkbox"/>	vsanScope

Step 3 Choose **enable** for each switch from the Feature Admin drop-down list.

Step 4 Repeat step 3 for all switches in the network.



Note A warning displays if you do not enable CFS for all switches in the network for this feature.

Step 5 Check the **Master** check box for the switch to act as the merge master for this feature.

Step 6 Choose **commit Changes** from the Config Action drop-down list for each switch that you enabled for CFS.

Step 7 Click the **Servers** tab in the Information pane.

You see the configuration for this feature based on the master switch (see [Figure 7-11](#)).

Step 8 Modify the feature configuration. For example, right-click the name in the Master column and choose **Create Row** to create a server for NTP.

- Enter the ID and the Name or IP Address of the NTP server.
- Set the **Mode** radio button and optionally check the **Preferred** check box.
- Click **Create** to add the server.

Figure 7-11 Servers Tab

Master	Id	Name or IP Address	IP Address Type	Preferred	Preferred
sw172-22-46-220	1	171.69.21.28	ipv4	peer	<input checked="" type="checkbox"/>
sw172-22-46-220	2	1.2.3.4	ipv4	peer	<input type="checkbox"/>
sw172-22-46-220	3	171.69.16.26	ipv4	server	<input checked="" type="checkbox"/>
sw172-22-46-220	4	171.69.16.23	ipv4	peer	<input checked="" type="checkbox"/>

Step 9 Click the **Delete Row** icon to delete a row.

If you make any changes, the status automatically changes to Pending (see [Figure 7-12](#)).

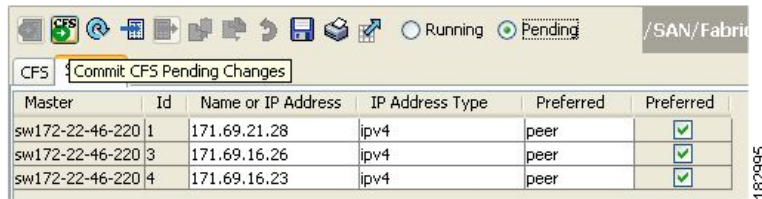
Figure 7-12 Status Change to Pending

Master	Id	Name or IP Address	IP Address Type	Preferred	Preferred
sw172-22-46-220	1	171.69.21.28	ipv4	peer	<input checked="" type="checkbox"/>
sw172-22-46-220	3	171.69.16.26	ipv4	server	<input checked="" type="checkbox"/>
sw172-22-46-220	4	171.69.16.23	ipv4	peer	<input checked="" type="checkbox"/>

Send comments to nx5000-docfeedback@cisco.com

Step 10 Click the **Commit CFS Pending Changes** icon to save the changes (see [Figure 7-13](#)).

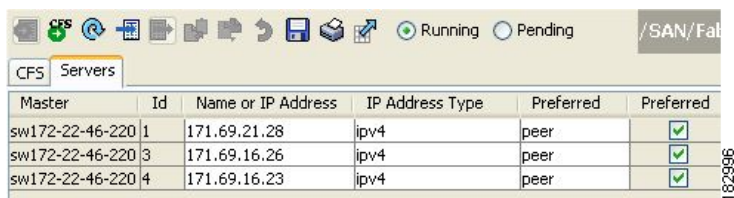
Figure 7-13 Commit CFS Pending Changes



Master	Id	Name or IP Address	IP Address Type	Preferred	Preferred
sw172-22-46-220	1	171.69.21.28	ipv4	peer	<input checked="" type="checkbox"/>
sw172-22-46-220	3	171.69.16.26	ipv4	peer	<input checked="" type="checkbox"/>
sw172-22-46-220	4	171.69.16.23	ipv4	peer	<input checked="" type="checkbox"/>

Step 11 The status changes to Running (see [Figure 7-14](#)).

Figure 7-14 Status Change to Running



Master	Id	Name or IP Address	IP Address Type	Preferred	Preferred
sw172-22-46-220	1	171.69.21.28	ipv4	peer	<input checked="" type="checkbox"/>
sw172-22-46-220	3	171.69.16.26	ipv4	peer	<input checked="" type="checkbox"/>
sw172-22-46-220	4	171.69.16.23	ipv4	peer	<input checked="" type="checkbox"/>

Step 12 Choose **abortChanges** from the Config Action drop-down list for each switch that you enabled for CFS (see [Figure 7-15](#)).

Figure 7-15 Commit Configuration Changes



Switch	Feature Admin	Feature Oper	Global State	Config Action	Last Command	Last Result	Lock Owner Switch	Lock Owner User Name	Merge Status	Master	Scope
sw172-22-46-220	noSelection	enabled	enable	noSelection	commitChanges	success			success	<input checked="" type="checkbox"/>	Fabric gNetwork
sw172-22-46-221	noSelection	enabled	enable	noSelection	commitChanges				success	<input type="checkbox"/>	Fabric gNetwork
sw172-22-46-174	noSelection	enabled	enable	noSelection	commitChanges				success	<input type="checkbox"/>	Fabric gNetwork



Note Fabric Manager does not change the status to pending if **enable** is selected, because the pending status does not apply until the first actual change is made.

Step 13 Click the **Apply Changes** icon to commit the configuration changes for that feature and distribute the changes through CFS.



Note When using CFS with features such as device alias, you must choose **commit** at the end of each configuration. If the session is locked, you must exit the feature by selecting **abort**.

Send comments to nx5000-docfeedback@cisco.com

To configure the master or seed switch for distribution for each feature using Fabric Manager, perform this task:

-
- Step 1** Choose the feature that needs a merge master for CFS.
For example, expand **Switches > Events**, and then choose **CallHome** from the Physical Attributes pane. The Information pane shows that feature including a CFS tab.
 - Step 2** Click the **CFS** tab to display the CFS state for each switch in the network for that feature.
 - Step 3** Check the **Master column** check box for the switch to act as the merge master for this feature.
 - Step 4** Click the **Apply Changes** icon to select this switch as master for future CFS distributions.
-

CFS Example Using Device Manager

This procedure is an example of what you see when you use Device Manager to configure a feature that uses CFS. For specific procedures for features that use CFS, refer to that feature's documentation.

To configure a feature that uses CFS using Device Manager, perform this task:

-
- Step 1** Open the dialog box for any CFS-capable feature.
Device Manager checks to see whether CFS is enabled. It also checks to see if there is a lock on the feature by checking for at least one entry in the Owner table. If CFS is enabled and there is a lock, Device Manager sets the status to pending for that feature. You see a dialog box displaying the lock information.
 - Step 2** Click **Continue** or **Cancel** when prompted. If you continue, Device Manager remembers the CFS status.
 - Step 3** Choose **Admin > CFS (Cisco Fabric Services)** to view the user name of the CFS lock holder.
 - Step 4** Click the locked feature and click **Details**.
 - Step 5** Click the **Owners** tab and look in the UserName column.



Note Device Manager does not monitor the status of the feature across the network until you click **Refresh**. If a user on another CFS-enabled switch attempts to configure the same feature, they do not see the pending status. However, their configuration changes are rejected by your switch.

- Step 6** If CFS is enabled and there is no lock, Device Manager sets the status to running for that feature.
You then see a dialog box for the feature. As soon as you perform a creation, deletion, or modification, Device Manager changes the status to pending and displays the updated information from the pending database.
 - Step 7** View the CFS table for a feature. Device Manager only changes the status to running when **commit**, **clear**, or **abort** is selected and applied. Device Manager will not change the status to pending if **enable** is selected, because the pending status does not apply until the first actual change is made.
The Last Command and Result fields are blank if the last command is **noOp**.
-

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

**Note**

When using CFS with features like device alias, you must choose **commit** at the end of each configuration. If the session is locked, you must exit the feature by selecting **abort**.

Default Settings

Table 7-1 lists the default settings for CFS configurations.

Table 7-1 *Default CFS Parameters*

Parameters	Default
CFS distribution on the switch	Enabled.
Database changes	Implicitly enabled with the first configuration change.
Application distribution	Differs based on application.
Commit	Explicit configuration is required.
CFS over IP	Disabled.
IPv4 multicast address	239.255.70.83.
IPv6 multicast address	ff15::eff:4653.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)



Configuring Ethernet Interfaces

This section describes the configuration of the Ethernet interfaces on the Cisco Nexus 5000 Series switches. It includes the following sections:

- [Information About Ethernet Interfaces, page 8-1](#)
- [Configuring Ethernet Interfaces, page 8-1](#)
- [Displaying Interface Information, page 8-1](#)

Information About Ethernet Interfaces

The Ethernet ports can operate as standard Ethernet interfaces connected to servers or to a LAN.

The Ethernet interfaces also support Fibre Channel over Ethernet (FCoE). FCoE allows the physical Ethernet link to carry both Ethernet and Fibre Channel traffic.

On the Cisco Nexus 5000 Series switch, the Ethernet interfaces are enabled by default.

Configuring Ethernet Interfaces

Fabric Manager and Device Manager display configuration settings and status information about the physical Ethernet interfaces on Cisco Nexus 5000 Series switches. However, you cannot change the configuration for physical Ethernet interfaces using Fabric Manager or Device Manager.

Displaying Interface Information

Fabric Manager and Device Manager display configuration settings and status information about the physical Ethernet interfaces on Cisco Nexus 5000 Series switches.

To display Ethernet interfaces using Fabric Manager, follow these steps:

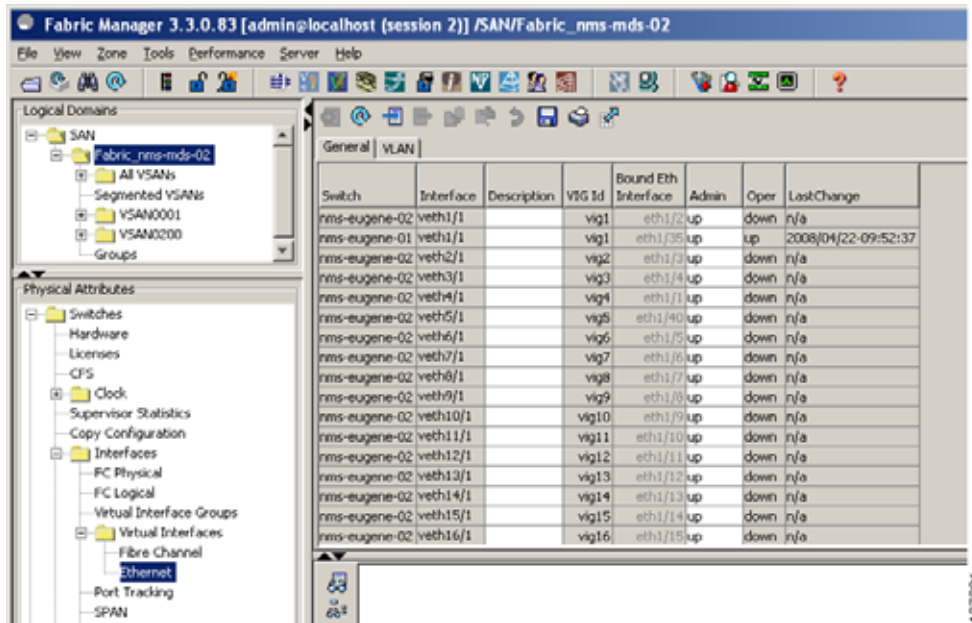
Step 1 In the Physical Attributes pane, expand **Switches > Interfaces** and then select **Ethernet**.

You see the Ethernet interface information pane as shown in [Figure 8-1](#).

The General tab displays the description, speed, MAC address and status for each interface.

Send comments to nx5000-docfeedback@cisco.com

Figure 8-1 Ethernet Information Pane



- Step 2 Select the **VLAN** tab to display the VLAN assigned to each interface.
- Step 3 Select the **CDP Neighbors** tab to display the CDP neighbor assigned to each interface.

Default Physical Ethernet Settings

The following table lists the default settings for all physical Ethernet interfaces:

Parameter	Default Setting
Debounce	Enable, 100 milliseconds
Duplex	Auto (full-duplex)
Encapsulation	ARPA
MTU ¹	1500 bytes
Port Mode	Access
Speed	Auto (10000)

1. MTU cannot be changed per-physical Ethernet interface. You modify MTU by selecting maps of QoS classes.



Configuring Virtual Interfaces

This section describes the configuration of virtual interfaces on the Cisco Nexus 5000 Series switches. It includes the following sections:

- [Information About Virtual Interfaces, page 9-1](#)
- [Configuring Virtual Interfaces, page 9-1](#)

Information About Virtual Interfaces

Cisco Nexus 5000 Series switches support I/O consolidation (IOC), which allows Fibre Channel and Ethernet traffic to be carried on the same physical Ethernet connection between the switch and the servers. For additional information about IOC, see [Chapter 1, “Product Overview.”](#)

The concept of virtual interface is used to emulate the logical connections that are carried on the same physical Ethernet. The Cisco Nexus 5000 Series switch supports virtual Ethernet and virtual Fibre Channel interfaces.

For configuration purposes, a virtual Ethernet or virtual Fibre Channel interface is implemented as a Layer 2 subinterface of the physical Ethernet interface. Logical features (such as VLAN and ACL) that can be configured on Ethernet interfaces can be configured on individual virtual Ethernet interfaces. Logical Fibre Channel features (such as VSAN) can be configured on virtual Fibre Channel interfaces.



Note

Virtual interfaces are created with the administrative state set to down. You need to explicitly configure the administrative state to bring the virtual interface into operation.

Configuring Virtual Interfaces

This section describes how to configure virtual interfaces, and it includes the following topics:

- [Creating a Virtual Interface Group, page 9-2](#)
- [Using the Virtual Interface Group Wizard, page 9-3](#)
- [Binding a VIG to a Physical Ethernet Interface, page 9-4](#)
- [Deleting a Virtual Interface Group, page 9-4](#)
- [Using the Virtual Interface Wizard, page 9-5](#)
- [Creating a Virtual Ethernet Interface, page 9-6](#)

Send comments to nx5000-docfeedback@cisco.com

- [Deleting a Virtual Ethernet Interface, page 9-7](#)
- [Creating a Virtual Fibre Channel Interface, page 9-7](#)
- [Deleting a Virtual Fibre Channel Interface, page 9-8](#)

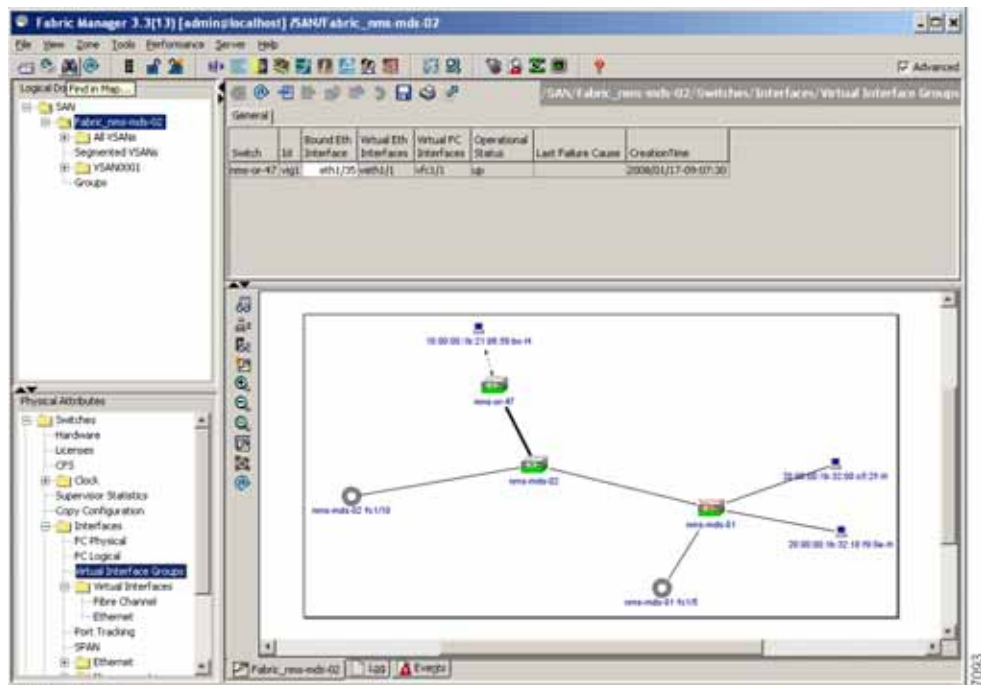
Creating a Virtual Interface Group

To create a virtual interface group (VIG), perform this task:

- Step 1** In the Physical Attributes pane, expand **Switches** > **Interfaces** > **Ethernet** > **FCoE**, and then choose **Virtual Interface Groups**.

You see the Virtual Interface Groups in the information pane as shown in [Figure 9-1](#).

Figure 9-1 Virtual Interface Group Information Pane




- Step 2** In the Virtual Interface Group information pane toolbar, click the **Create Row** icon.

You see the Create VIG dialog box as shown in [Figure 9-2](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 9-2 Create VIG Dialog Box



- Step 3** Choose a switch in the Switch pull-down menu.
- Step 4** Enter the Virtual Interface Group number in the Id field.
- Step 5** (Optional) Select a physical Ethernet interface to bind to the Virtual Interface Group.
The newly created virtual interface group is displayed in the Virtual Interface Groups table. In the Create Virtual Interface Group dialog box, the Id field increments by 1.
- Step 6** (Optional) Repeat steps 3 through 6 to create additional virtual interface groups.
- Step 7** Click **Close** to finish.
-

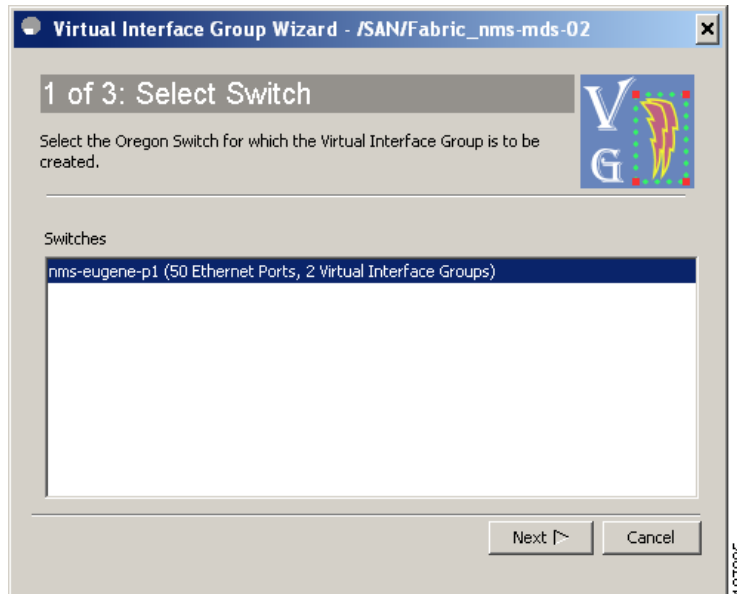
Using the Virtual Interface Group Wizard

To create a virtual interface group using the Virtual Interface Group wizard, perform this task:

- Step 1** Click the **Virtual Interface Group** button on the tool bar.
You see the Virtual Interface Group wizard dialog box as shown in [Figure 9-3](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 9-3 Virtual Interface Group Wizard



- Step 2** Follow the directions in the Virtual Interface Group Wizard. Choose the switch where the Virtual Interface Group will be created, the physical Ethernet port associated with the Virtual Interface Group, and the Virtual Interface Group ID.
- Step 3** Click **Finish** to commit and distribute the change.

Binding a VIG to a Physical Ethernet Interface

To bind the VIG to a physical Ethernet port, perform this task:

- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces > Ethernet > FCoE**, and then choose **Virtual Interface Groups**.
- You see the Virtual Interface Groups in the information pane as shown in [Figure 9-1](#).
- Step 2** Click the **Bound Eth Interface** entry for the Virtual Interface Group that needs to be modified. Enter the physical interface.
- Step 3** Click **Apply Changes** to commit and distribute the change.

Deleting a Virtual Interface Group

To delete a VIG, perform this task:

- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces > Ethernet > FCoE**, and then choose **Virtual Interface Groups**.

Send comments to nx5000-docfeedback@cisco.com

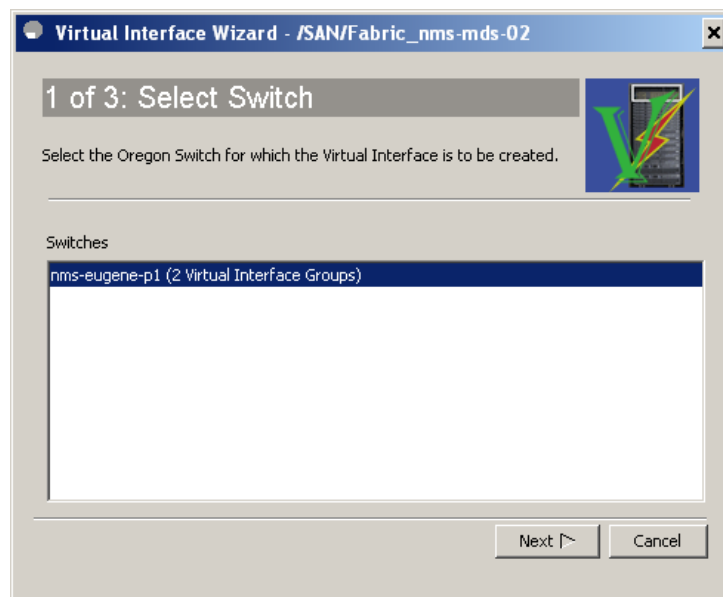
You see the Virtual Interface Groups in the information pane as shown in [Figure 9-1](#).

- Step 2** Select the Virtual Interface Group to be deleted.
- Step 3** Click the **Delete Row** icon.
You see a dialog box asking you to confirm the deletion.
- Step 4** Click **Yes** to confirm the deletion.
- Step 5** Click the **Apply** icon to apply the change.

Using the Virtual Interface Wizard

- Step 1** Click the **Virtual Interface** button on the tool bar.
You see the Virtual Interface wizard dialog box as shown in [Figure 9-4](#).

Figure 9-4 Virtual Interface Wizard



- Step 2** Follow the directions in the Virtual Interface Wizard. Select the switch where the virtual interface will be created, and the virtual interface group to contain the virtual interface, and specify the virtual interface type (Ethernet or Fibre Channel). The wizard sets the virtual interface ID to 1.
- Step 3** Click **Finish** to commit and distribute the change.

Send comments to nx5000-docfeedback@cisco.com

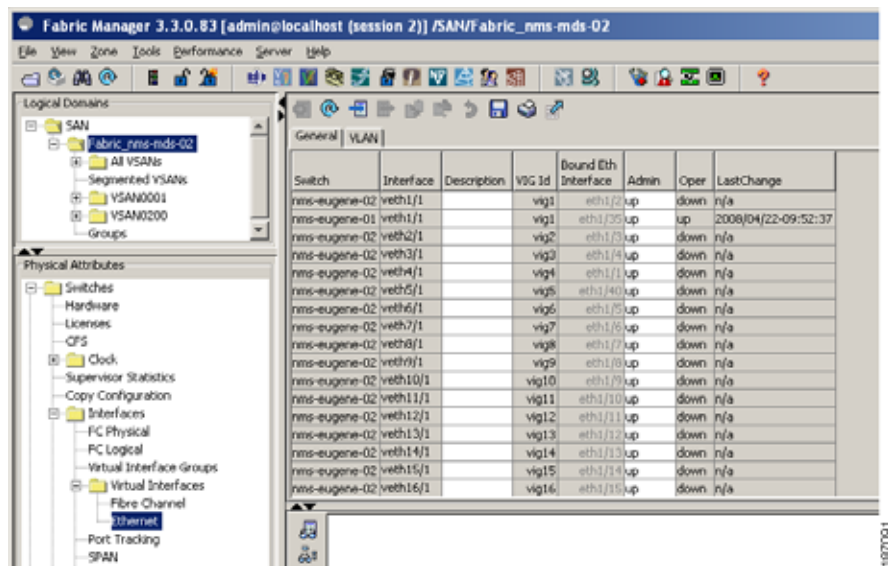
Creating a Virtual Ethernet Interface

To create a virtual Ethernet interface, perform this task:

- Step 1** In the Physical Attributes pane, expand **Switches**, expand **Interfaces**, expand **Ethernet**, expand **FCoE**, expand **Virtual Interfaces**, and then choose **Ethernet**.

You see the Virtual Ethernet information pane as shown in [Figure 9-5](#).


Figure 9-5 Virtual Ethernet



- Step 2** In the Information pane toolbar, click the **Create Row** icon.

You see the Virtual Ethernet dialog box as shown in [Figure 9-6](#).

Figure 9-6 Create Virtual Ethernet



- Step 3** Select a switch in the Switch pull-down menu.

- Step 4** Enter the virtual interface group ID in the VIG ID field. The virtual interface ID is set to 1.

- Step 5** Click **Create**.

Send comments to nx5000-docfeedback@cisco.com

Step 6 Click the Apply icon to apply the change.

Deleting a Virtual Ethernet Interface

To delete a virtual Ethernet interface, perform this task:

Step 1 In the Physical Attributes pane, expand **Switches > Interfaces > Ethernet > FCoE**, and then choose **Virtual Ethernet Interfaces**.

You see the Virtual Ethernet information pane as shown in [Figure 9-5](#).

Step 2 Select a virtual Ethernet row from the information pane.

Step 3 In the Information pane toolbar, click the **Delete Row** icon.

Step 4 Confirm the deletion in the dialog box.

Creating a Virtual Fibre Channel Interface


To create a virtual Fibre Channel interface, perform this task:

Step 1 In the Physical Attributes pane, expand **Switches > Interfaces > Ethernet > FCoE**, and then choose **Virtual FC Interfaces**.

Step 2 In the Information pane toolbar, click the **Create Row** icon.

You see the Virtual Ethernet dialog box as shown in [Figure 9-7](#).

Figure 9-7 Create Virtual Fibre Channel



Step 3 Select a switch in the Switch pull-down menu.

Step 4 Enter the Virtual Interface Group ID in the VIG ID field. The interface ID is set to a value of 1.

Step 5 Click **Create**.

Send comments to nx5000-docfeedback@cisco.com

Deleting a Virtual Fibre Channel Interface

To delete a virtual Fibre Channel interface, perform this task:

Step 1 In the Physical Attributes pane, expand **Switches > Interfaces > Ethernet > FCoE**, and then choose **Virtual FC Interfaces**.

You see the Virtual Fibre Channel information pane.

Step 2 Select a virtual Fibre Channel row from the information pane.

Step 3 In the Information pane toolbar, click the **Delete Row** icon.

Step 4 Confirm the deletion in the dialog box.



Configuring Fibre Channel Interfaces

This chapter describes interface configuration for Fibre Channel interfaces and virtual Fibre Channel interfaces. This chapter includes the following sections:

- [Information About Fibre Channel Interfaces, page 10-1](#)
- [Configuring Fibre Channel Interfaces, page 10-8](#)
- [Verifying Fibre Channel Interfaces, page 10-12](#)
- [Default Settings, page 10-14](#)

Information About Fibre Channel Interfaces

This section describes Fibre Channel interfaces and virtual Fibre Channel interfaces. This section includes the following topics:

- [Licensing Requirements, page 10-1](#)
- [Physical Fibre Channel Interfaces, page 10-1](#)
- [Virtual Fibre Channel Interfaces, page 10-2](#)
- [Interface Modes, page 10-2](#)
- [Interface States, page 10-5](#)
- [Buffer-to-Buffer Credits, page 10-7](#)

Licensing Requirements

On Cisco Nexus 5000 Series switches, Fibre Channel capability is included in the Storage Protocol Services License.

Ensure that you have the correct license installed (N5010SS or N5020SS) before using Fibre Channel interfaces and capabilities.

Physical Fibre Channel Interfaces

Cisco Nexus 5000 Series switches provide up to eight physical Fibre Channel uplinks. The Fibre Channel interfaces are supported on optional expansion modules. The Fibre Channel plus Ethernet expansion module contains four Fibre Channel interfaces.

Send comments to nx5000-docfeedback@cisco.com

Each Fibre Channel port can be used as a downlink (connected to a server) or as an uplink (connected to the data center SAN network). The Fibre Channel interfaces support the following modes: F, NP, E, TE, and SD.

Virtual Fibre Channel Interfaces

Fibre Channel over Ethernet (FCoE) encapsulation allows a physical Ethernet cable to simultaneously carry Fibre Channel and classic Ethernet (CE) traffic. In the Cisco Nexus 5000 Series switches, an FCoE-capable physical Ethernet interface can carry traffic for one logical CE interface and one logical Fibre Channel interface. The logical interfaces are configured in the Cisco Nexus 5000 Series switch as virtual interfaces. A virtual Fibre Channel interface represents the logical Fibre Channel interface.

A virtual Fibre Channel is configured as a subinterface of a virtual interface group (VIG).

Virtual Fibre Channel interfaces support only F mode, and offer a subset of the features that are supported on physical Fibre Channel interfaces.

The following capabilities are not supported for virtual Fibre Channel interfaces:

- SAN port channels.
- VSAN trunking. The virtual Fibre Channel is associated with one VSAN.
- The SPAN destination cannot be a virtual Fibre Channel interface.
- Buffer-to-buffer credits.
- Exchange link parameters (ELP), or Fabric Shortest Path First (FSPF) protocol.
- Configuration of physical attributes (speed, rate, mode, transmitter information, MTU size).
- Port tracking.

Interface Modes

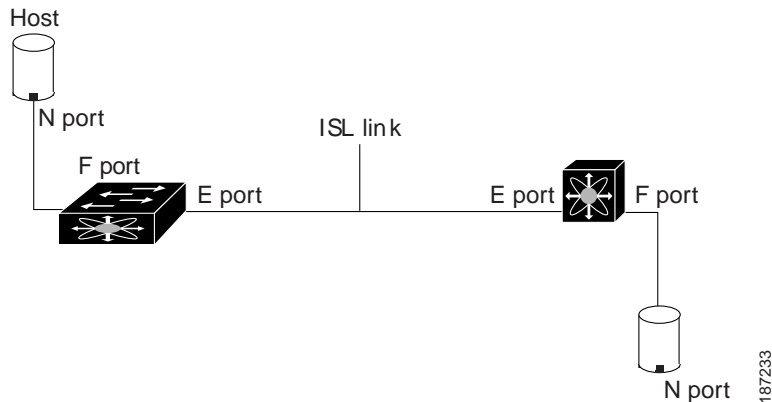
Each physical Fibre Channel interface in a switch may operate in one of several port modes: E mode, TE mode, F mode, and SD mode (see [Figure 10-1](#)). A physical Fibre Channel interface can be configured as an E port, an F port, or an SD port. Interfaces may also be configured in Auto mode; the port type is determined during interface initialization.

In NPV mode, Fibre Channel interfaces may operate in NP mode, F mode or SD mode. For additional information about NPV mode, see [Chapter 12, “Configuring N-Port Virtualization.”](#)

Virtual Fibre Channel interfaces can only be configured in F mode.

Send comments to nx5000-docfeedback@cisco.com

Figure 10-1 Switch Port Modes



Note

Interfaces are automatically assigned VSAN 1 by default. See [Chapter 15, “Configuring and Managing VSANs.”](#)

Each interface has an associated administrative configuration and an operational status:

- The administrative configuration does not change unless you modify it. This configuration has various attributes that you can configure in administrative mode.
- The operational status represents the current status of a specified attribute such as 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).

The following sections provide a brief description of each interface mode:

- [E Port, page 10-3](#)
- [F Port, page 10-4](#)
- [NP Port, page 10-4](#)
- [TE Port, page 10-4](#)
- [SD Port, page 10-4](#)
- [Auto Mode, page 10-4](#)

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. E ports support class 3 and class F service.

An E port connected to another switch may also be configured to form a SAN port channel (see [Chapter 14, “Configuring SAN Port Channels”](#)).

Send comments to nx5000-docfeedback@cisco.com

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 3 service.

NP Port

An NP port is a port on a device that is in NPV mode and connected to the core NPV switch through an F port. NP ports operate like N ports that function as proxies for multiple physical N ports.

For more details about NP ports and NPV, see [Chapter 12, “Configuring N-Port Virtualization.”](#)

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 connect to another Cisco Nexus 5000 Series switch or a Cisco MDS 9000 Family switch. They expand the functionality of E ports to support the following:

- VSAN trunking
- 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 VSAN trunking in the Cisco Nexus 5000 Series (see [Chapter 13, “Configuring VSAN Trunking”](#)). TE ports support 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 monitors network traffic that passes through 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, instead they 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.

Auto Mode

Interfaces configured in auto mode can operate in one of the following modes: F 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 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 Nexus 5000 Series or Cisco MDS 9000 Family, it may become operational in TE port mode (see [Chapter 13, “Configuring VSAN Trunking”](#)).

SD ports are not determined during initialization and are administratively configured.

Send comments to nx5000-docfeedback@cisco.com

Interface States

The interface state depends on the administrative configuration of the interface and the dynamic state of the physical link. The following sections describe the states and configuration that influence the state:

- [Administrative States, page 10-5](#)
- [Operational States, page 10-5](#)
- [Reason Codes, page 10-5](#)

Administrative States

The administrative state refers to the administrative configuration of the interface. [Table 10-1](#) describes the administrative states.

Table 10-1 Administrative States

Administrative State	Description
Up	Interface is enabled.
Down	Interface is disabled. If you administratively disable an interface by shutting down that interface, the physical link layer state change is ignored.

Operational States

The operational state indicates the current operational state of the interface. [Table 10-2](#) describes the operational states.

Table 10-2 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 10-3](#) describes the reason codes for operational states.

Table 10-3 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 Table 10-4 .

Send comments to nx5000-docfeedback@cisco.com

**Note**

Only some of the reason codes are listed in [Table 10-4](#).

If the administrative state is up and the operational state is down, the reason code differs based on the nonoperational reason code. [Table 10-4](#) describes the reason codes for nonoperational states.

Table 10-4 Reason Codes for Nonoperational States

Reason Code (long version)	Description	Applicable Modes
Link failure or not connected	The physical layer link is not operational.	All
SFP not present	The small form-factor pluggable (SFP) hardware is not plugged in.	
Initializing	The physical layer link is operational and the protocol initialization is in progress.	All
Reconfigure fabric in progress	The fabric is currently being reconfigured.	
Offline	The switch software waits for the specified R_A_TOV time before retrying initialization.	
Inactive	The interface VSAN is deleted or is in a suspended state. To make the interface operational, assign that port to a configured and active VSAN.	
Hardware failure	A hardware failure is detected.	
Error disabled	Error conditions require administrative attention. Interfaces may be error-disabled for various reasons. For example: <ul style="list-style-type: none"> • Configuration failure. • Incompatible buffer-to-buffer credit configuration. To make the interface operational, you must first fix the error conditions causing this state and then administratively shut down or enable the interface.	
Isolation because limit of active port channels is exceeded.	The interface is isolated because the switch is already configured with the maximum number of active SAN port channels.	

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Table 10-4 Reason Codes for Nonoperational States (continued)

Reason Code (long version)	Description	Applicable Modes
Isolation due to ELP failure	The port negotiation failed.	Only E ports and TE ports
Isolation due to ESC failure	The port negotiation failed.	
Isolation due to domain overlap	The Fibre Channel domains (fcdomain) overlap.	
Isolation due to domain ID assignment failure	The assigned domain ID is not valid.	
Isolation due to the other side of the link E port isolated	The E port at the other end of the link is isolated.	
Isolation due to invalid fabric reconfiguration	The port is isolated due to fabric reconfiguration.	
Isolation due to domain manager disabled	The fcdomain feature is disabled.	
Isolation due to zone merge failure	The zone merge operation failed.	
Isolation due to VSAN mismatch	The VSANs at both ends of an ISL are different.	
port channel administratively down	The interfaces belonging to the SAN port channel are down.	Only SAN port channel interfaces
Suspended due to incompatible speed	The interfaces belonging to the SAN port channel have incompatible speeds.	
Suspended due to incompatible mode	The interfaces belonging to the SAN port channel have incompatible modes.	
Suspended due to incompatible remote switch WWN	An improper connection is detected. All interfaces in a SAN port channel must be connected to the same pair of switches.	

Buffer-to-Buffer Credits

Buffer-to-buffer credits (BB_credits) are a flow-control mechanism to ensure that Fibre Channel interfaces do not drop frames. BB_credits are negotiated on a per-hop basis.

In Cisco Nexus 5000 Series switches, the BB_credit mechanism is used on Fibre Channel interfaces but not on virtual Fibre Channel interfaces. Virtual Fibre Channel interfaces provide flow control based on capabilities of the underlying physical Ethernet interface.

The receive BB_credit value (fcrxbbcredit) may be configured for each Fibre Channel interface. In most cases, you do not need to modify the default configuration.



Note

The receive BB_credit values depend on the port mode, as follows:

- For physical Fibre Channel interfaces, the default value is 16 for F mode and E mode interfaces. This value can be changed as required. The maximum value is 64.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

- For virtual Fibre Channel interfaces, BB_credits are not used.

Configuring Fibre Channel Interfaces

This section describes how to configure Fibre Channel interfaces, and includes the following topics:

- [Configuring a Fibre Channel Interface, page 10-8](#)
- [Setting the Interface Administrative State, page 10-9](#)
- [Configuring Interface Modes, page 10-9](#)
- [Configuring the Interface Description, page 10-9](#)
- [Configuring Administrative Speeds, page 10-10](#)
- [Configuring SD Port Frame Encapsulation, page 10-10](#)
- [Configuring Receive Data Field Size, page 10-11](#)
- [Understanding Bit Error Thresholds, page 10-11](#)
- [Configuring Buffer-to-Buffer Credits, page 10-12](#)

Configuring a Fibre Channel Interface

You can configure native Fibre Channel interfaces using Fabric Manager by expanding **Switches > Interfaces > FC Physical** from the Physical Attributes pane.

Figure 10-2 shows an example of the Information pane for Fibre Channel Interfaces.

Figure 10-2 Native Fibre Channel Interface Configuration

Switch	Interface	Mode Admin	Mode Oper	Port VSAN	Dynamic VSAN	Description	Speed Admin	Speed Oper	Rate Mode	Status Service	Status Admin	Status Oper	FailureCause	Was Enabled	LastChange
nms-nds-02	fc1/11	auto	auto	1	n/a		auto	n/a	dedicated	in	up	down	linkFailure	true	2008/03/12-10:40:43
nms-nds-02	fc1/12	auto	auto	1	n/a		auto	n/a	dedicated	in	up	down	linkFailure	false	n/a
nms-nds-01	fc1/1	auto	F	888	n/a		auto	2 Gb	dedicated	in	up	up	none	true	2008/02/21-14:04:35
nms-nds-02	fc1/3	auto	auto	1	n/a		auto	n/a	dedicated	in	up	down	linkFailure	true	2008/03/12-14:30:23
nms-nds-02	fc1/4	auto	auto	1	n/a		auto	n/a	dedicated	in	up	down	linkFailure	false	n/a
nms-nds-01	fc1/2	auto	FL	1	n/a		auto	3 Gb	dedicated	in	up	up	none	true	2008/02/29-13:45:49
nms-nds-02	fc1/5	auto	auto	1	n/a		auto	n/a	dedicated	in	up	down	linkFailure	false	n/a
nms-nds-02	fc1/6	auto	auto	1	n/a		auto	n/a	dedicated	in	up	down	linkFailure	false	n/a
nms-nds-01	fc1/3	auto	F	1	n/a		auto	2 Gb	dedicated	in	up	up	none	true	2008/03/08-13:43:49
nms-nds-02	fc1/7	auto	auto	1	n/a		auto	n/a	dedicated	in	up	down	linkFailure	false	n/a
nms-nds-02	fc1/8	auto	auto	1	n/a		auto	n/a	dedicated	in	up	down	linkFailure	false	n/a
nms-nds-01	fc1/4	auto	auto	1	n/a		auto	n/a	dedicated	in	up	down	linkFailure	false	n/a
nms-nds-01	fc1/9	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	linkFailure	false	n/a
nms-nds-02	fc1/9	auto	auto	1	n/a		auto	n/a	dedicated	in	up	down	linkFailure	false	n/a
nms-nds-02	fc1/10	auto	TE	1	n/a		auto	2 Gb	dedicated	in	up	up	none	true	2008/03/11-09:52:29
nms-nds-02	fc1/11	auto	auto	1	n/a		auto	n/a	dedicated	in	up	down	linkFailure	false	n/a
nms-nds-01	fc1/5	E	TE				auto	2 Gb	dedicated	in	up	up	none	true	2008/03/11-09:52:45
nms-nds-02	fc1/12	auto	FL	888	n/a		auto	3 Gb	dedicated	in	up	up	none	true	2008/03/03-15:16:12
nms-nds-01	fc1/11	auto	auto	1	n/a		auto	n/a	dedicated	in	down	down	linkFailure	false	n/a

You can configure virtual Fibre Channel interfaces using Fabric Manager by expanding **Switches > Interfaces > Ethernet > FCoE > Virtual FC Interfaces** from the Physical Attributes pane.

Figure 10-3 shows an example of the Information pane for virtual Fibre Channel Interfaces.

Send comments to nx5000-docfeedback@cisco.com

Figure 10-3 Virtual Fibre Channel Interface Configuration

Switch	Interface	Description	VIG Id	Bound Eth Interface	Port VSAN	Mode Admin	Mode Oper	Status Service	Status Admin	Status Oper	FailureCause	Was Enabled	LastChange
nms-eugene-p1	vfc6/1		vig6	0	1 F	auto	in	down	down	adminDown	False	n/a	
nms-eugene-p1	vfc27/1		vig27	eth1/27	1 F	auto	in	up	down	none	False	n/a	

Setting the Interface Administrative State

To disable or enable an interface using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**. For a virtual Fibre Channel Interface, expand **Switches > Interfaces > Ethernet > FCoE > Virtual FC Interfaces**.
- You see the interface configuration in the Information pane.
- Step 2** Click the **General** tab.
- Step 3** Click **Status Admin**.
- You see a drop-down box with a choice of up or down.
- Step 4** Set the status to down (disable) or up (enable).
- Step 5** Click **Apply Changes**.
-

Configuring Interface Modes

To configure the interface mode using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**. For a virtual Fibre Channel Interface, expand **Switches > Interfaces > Ethernet > FCoE > Virtual FC Interfaces**.
- You see the interface configuration in the Information pane.
- Step 2** Click the **General** tab.
- Step 3** Click **Mode Admin**. Choose the desired mode from the pull-down list.
- Step 4** Click **Apply Changes** icon.

Configuring the Interface Description

Interface descriptions should help you identify the traffic or use for that interface. The interface description can be any alphanumeric string.

Send comments to nx5000-docfeedback@cisco.com

To configure the interface using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**. For a virtual Fibre Channel Interface, expand **Switches > Interfaces > Ethernet > FCoE > Virtual FC Interfaces**.
You see the interface configuration in the Information pane.
 - Step 2** Click the **General** tab.
 - Step 3** Click **Description**. Enter the desired text.
 - Step 4** (Optional) Set additional configuration parameters using the other tabs.
 - Step 5** Click **Apply Changes** icon.

Configuring Administrative Speeds

Administrative speed can be configured on a physical Fibre Channel interface (but not on a virtual Fibre Channel interface). By default, the administrative speed for an interface is automatically calculated by the switch.



Caution

Changing the administrative speed is a disruptive operation.

To configure administrative speed of the interface using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**.
You see the interface configuration in the Information pane.
 - Step 2** Click the **General** tab.
 - Step 3** Click **Speed Admin**. Set the desired speed from the drop-down list.
The number indicates the speed in megabits per second (Mbps). You can set the speed to 1-Gbps, 2-Gbps, 4-Gbps, or **auto** (default).
 - Step 4** Click **Apply Changes**.
-

Autosensing

Autosensing speed is enabled on all 4-Gbps interfaces by default. This configuration enables the interfaces to operate at speeds of 1 Gbps, 2 Gbps, or 4 Gbps on the 4-Gbps ports. When autosensing is enabled for an interface operating in dedicated rate mode, 4-Gbps of bandwidth is reserved, even if the port negotiates at an operating speed of 1-Gbps or 2-Gbps.

Configuring SD Port Frame Encapsulation

You can set the frame format to EISL for all frames transmitted by the interface in SD port mode. If you set the frame encapsulation to EISL, all outgoing frames are transmitted in the EISL frame format for all SPAN sources.

Send comments to nx5000-docfeedback@cisco.com

See the *Cisco Cisco Nexus 5000 Series Family CLI Configuration Guide* to configure frame encapsulation on an interface.

Configuring Receive Data Field Size

You can configure the receive data field size for native Fibre Channel interfaces (but not for virtual Fibre Channel interfaces). If the default data field size is 2112 bytes, the frame length will be 2148 bytes.

To configure the receive data field size using Fabric Manager, perform this task:

- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**. You see the interface configuration in the Information pane.
- Step 2** Click the **Other** tab and set the RxDataFieldSize field (see [Figure 10-4](#)).

Figure 10-4 Changing Rx Data Size

The screenshot shows the Fabric Manager interface with the 'Other' tab selected in the Physical Attributes pane. The table below represents the data visible in the interface.

Switch	Interface	PortChannelId	Auto Port Channel	Fabric WWN	Mtu	RxDataFieldSize	HoldTime
nms-mds-02	Fc1/1	channel1	<input type="checkbox"/>	20-01:00:0d:ec:0d:d0:00	2112	2112	0
nms-mds-02	Fc1/2	none	<input type="checkbox"/>	20-02:00:0d:ec:0d:d0:00	2112	2112	0
nms-mds-01	Fc1/1	none	<input type="checkbox"/>	20-01:00:0d:ec:4e:87:40	2112	2112	0
nms-mds-02	Fc1/3	none	<input type="checkbox"/>	20-03:00:0d:ec:0d:d0:00	2112	2112	0
nms-mds-02	Fc1/4	none	<input type="checkbox"/>	20-04:00:0d:ec:0d:d0:00	2112	2112	0
nms-mds-01	Fc1/2	none	<input type="checkbox"/>	20-02:00:0d:ec:4e:87:40	2112	2112	0
nms-mds-02	Fc1/5	none	<input type="checkbox"/>	20-05:00:0d:ec:0d:d0:00	2112	2112	0
nms-mds-02	Fc1/6	none	<input type="checkbox"/>	20-06:00:0d:ec:0d:d0:00	2112	2112	0
nms-mds-01	Fc1/3	none	<input type="checkbox"/>	20-03:00:0d:ec:4e:87:40	2112	2112	0
nms-mds-02	Fc1/7	none	<input type="checkbox"/>	20-07:00:0d:ec:0d:d0:00	2112	2112	0
nms-mds-02	Fc1/8	none	<input type="checkbox"/>	20-08:00:0d:ec:0d:d0:00	2112	2112	0
nms-mds-01	Fc1/4	none	<input type="checkbox"/>	20-04:00:0d:ec:4e:87:40	2112	2112	0
nms-mds-02	Fc1/9	none	<input type="checkbox"/>	20-09:00:0d:ec:0d:d0:00	2112	2112	0
nms-mds-02	Fc1/10	none	<input type="checkbox"/>	20-0a:00:0d:ec:0d:d0:00	2112	2112	0
nms-mds-01	Fc1/5	none	<input type="checkbox"/>	20-05:00:0d:ec:4e:87:40	2112	2112	0
nms-mds-02	Fc1/11	none	<input type="checkbox"/>	20-0b:00:0d:ec:0d:d0:00	2112	2112	0
nms-mds-02	Fc1/12	none	<input type="checkbox"/>	20-0c:00:0d:ec:0d:d0:00	2112	2112	0
nms-mds-02	Fc1/13	none	<input type="checkbox"/>	20-0d:00:0d:ec:0d:d0:00	2112	2112	0
nms-mds-01	Fc1/6	none	<input type="checkbox"/>	20-06:00:0d:ec:4e:87:40	2112	2112	0

- Step 3** Click **Apply Changes**.

Understanding Bit Error Thresholds

The bit error rate threshold is used by the switch to detect an increased error rate before performance degradation seriously affects traffic.

The bit errors can occur for the following reasons:

- Faulty or bad cable.
- Faulty or bad GBIC or SFP.
- GBIC or SFP is specified to operate at 1 Gbps but is used at 2 Gbps.
- GBIC or SFP is specified to operate at 2 Gbps but is used at 4 Gbps.
- Short haul cable is used for long haul or long haul cable is used for short haul.

Send comments to nx5000-docfeedback@cisco.com

- Momentary synchronization loss.
- Loose cable connection at one or both ends.
- Improper GBIC or SFP connection at one or both ends.

A bit error rate threshold is detected when 15 error bursts occur in a 5-minute period. By default, the switch disables the interface when the threshold is reached. You can reenabling the interface.

You can configure the switch to not disable an interface when the threshold is crossed.

See the *Cisco Cisco Nexus 5000 Series CLI Configuration Guide* to disable the bit error threshold for an interface.



Note

The switch generates a syslog message when bit error threshold events are detected, even if the interface is configured not to be disabled by bit-error threshold events.

Configuring Buffer-to-Buffer Credits

The BB_credit scheme is not used for virtual Fibre Channel interfaces. To configure BB_credits for a native Fibre Channel interface using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**.
You see the interface configuration in the Information pane.
- Step 2** Choose the **Bb Credit** tab.
You see the buffer credits.
- Step 3** Set any of the buffer-to-buffer credits for an interface.
- Step 4** Click **Apply Changes**.
-

Verifying Fibre Channel Interfaces

The following topics describe the commands for displaying Fibre Channel interfaces:

- [Verifying SFP Transmitter Types, page 10-13](#)
- [Obtaining Interface Statistics, page 10-13](#)

Send comments to nx5000-docfeedback@cisco.com

Verifying SFP Transmitter Types

The SFP transmitter type can be displayed for a physical Fibre Channel interface (but not for a virtual Fibre Channel).

The small form-factor pluggable (SFP) hardware transmitters are identified by their acronyms when displayed.

To display the SFP types for an interface using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**.
You see the interface configuration in the Information pane.
- Step 2** Click the **Physical** tab to see the transmitter type for the selected interface.
-

Obtaining Interface Statistics

You can use Fabric Manager or Device Manager to collect interface statistics on any switch. These statistics are collected at intervals that you can set.




Note

In Fabric Manager, you can collect interface statistics by expanding **Switches > ISLs** and selecting **Statistics** from the Physical Attributes pane.

To obtain and display interface counters using Device Manager, perform this task:

-
- Step 1** Right-click an interface and choose **Monitor** in the Interface menu and choose **Ethernet Enabled** or **FC Enabled**.
You see the Interface Monitor dialog box.
- Step 2** Set both the number of seconds at which you want to poll the interface statistics and how you want the data represented in the Interval drop-down menus. For example, click **10s** and **LastValue/sec**.
- Step 3** Click any tab shown in [Figure 10-5](#) to view those related statistics.

Figure 10-5 Device Manager FC Interface Monitor Dialog Box



- Step 4** (Optional) Click the **Pencil** icon to reset the cumulative counters.
- Step 5** (Optional) Click the **Save** icon to save the gathered statistics to a file or click the **Print** icon to print the statistics.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Step 6 Click **Close** when you are finished gathering and displaying statistics.

Default Settings

Table 10-6 lists the default settings for native Fibre Channel interface parameters.

Table 10-5 *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

Table 10-6 lists the default settings for virtual Fibre Channel interface parameters.

Table 10-6 *Default Virtual Fibre Channel Interface Parameters*

Parameters	Default
Interface mode	Auto
Interface speed	n/a
Administrative state	Shutdown (unless changed during initial setup)
Trunk mode	n/a
Trunk-allowed VSANs	n/a
Interface VSAN	Default VSAN (1)
EISL encapsulation	n/a
Data field size	n/a



Configuring Domain Parameters

The Fibre Channel domain (fcdomain) feature performs principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions as described in the FC-SW-2 standards. The domains are configured on a per VSAN basis. If you do not configure a domain ID, the local switch uses a random ID.



Caution

Changes to fcdomain parameters should not be performed on a daily basis. These changes should be made by an administrator or individual who is completely familiar with switch operations.



Tip

When you change the configuration, be sure to save the running configuration. The next time you reboot the switch, the saved configuration is used. If you do not save the configuration, the previously saved startup configuration is used.

This chapter includes the following sections:

- [Information About Fibre Channel Domains, page 11-1](#)
- [Domain IDs, page 11-8](#)
- [FC IDs, page 11-15](#)
- [Displaying fcdomain Statistics, page 11-20](#)
- [Default Settings, page 11-21](#)

Information About Fibre Channel Domains

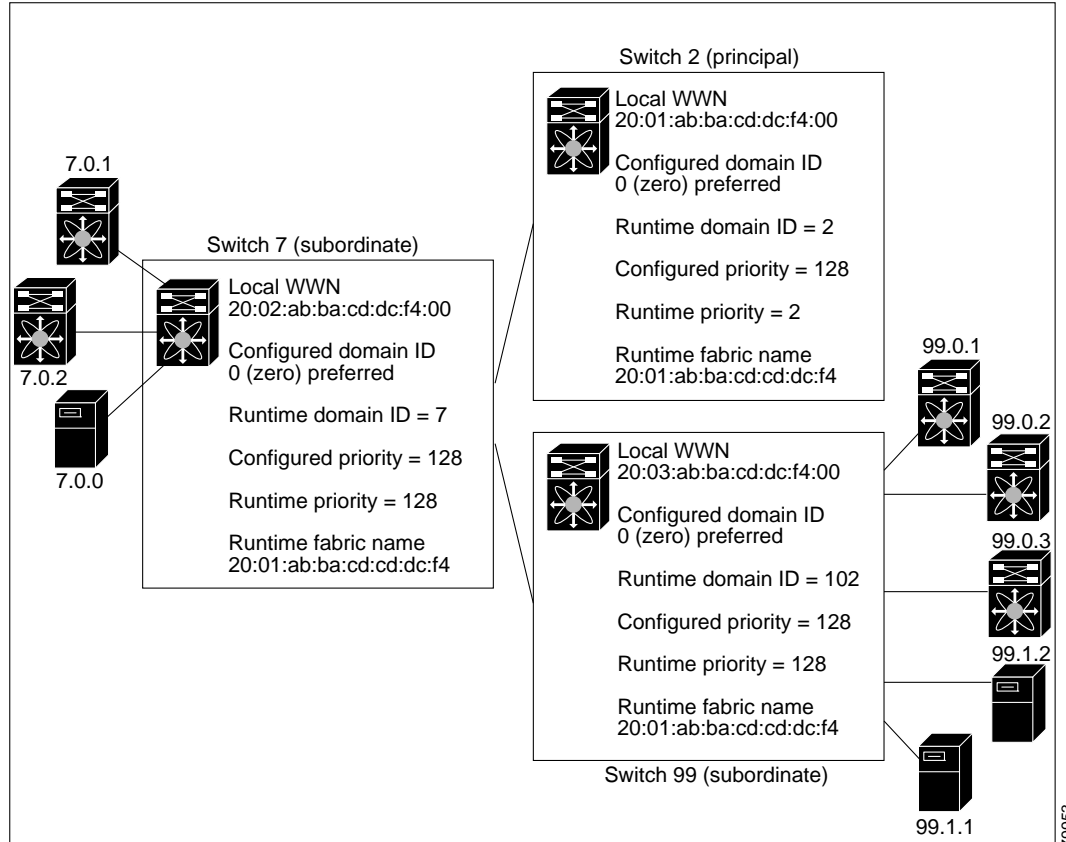
This section describes each fcdomain phase:

- **Principal switch selection**—This phase guarantees the selection of a unique principal switch across the fabric.
- **Domain ID distribution**—This phase guarantees each switch in the fabric obtains a unique domain ID.
- **FC ID allocation**—This phase guarantees a unique FC ID assignment to each device attached to the corresponding switch in the fabric.
- **Fabric reconfiguration**—This phase guarantees a resynchronization of all switches in the fabric to ensure they simultaneously restart a new principal switch selection phase.

See [Figure 11-1](#) for an example fcdomain configuration.

Send comments to nx5000-docfeedback@cisco.com

Figure 11-1 Sample fcdomain Configuration



Note

Domain IDs and VSAN values used in all procedures are only provided as examples. Be sure to use IDs and values that apply to your configuration.

This section describes the fcdomain feature and includes the following topics:

- [About Domain Restart, page 11-3](#)
- [Restarting a Domain, page 11-3](#)
- [About Switch Priority, page 11-4](#)
- [Configuring Switch Priority, page 11-4](#)
- [About fcdomain Initiation, page 11-5](#)
- [Enabling or Disabling fcdomains, page 11-5](#)
- [Setting Fabric Names, page 11-6](#)
- [About Incoming RCFs, page 11-6](#)
- [Rejecting Incoming RCFs, page 11-6](#)
- [About Autoreconfiguring Merged Fabrics, page 11-7](#)
- [Enabling Autoreconfiguration, page 11-7](#)

Send comments to nx5000-docfeedback@cisco.com

About Domain Restart

Fibre Channel domains can be started disruptively or nondisruptively. If you perform a disruptive restart, reconfigure fabric (RCF) frames are sent to other switches in the fabric and data traffic is disrupted on all the switches in the VSAN (including remotely segmented ISLs). If you perform a nondisruptive restart, build fabric (BF) frames are sent to other switches in the fabric and data traffic is disrupted only on the switch.

If you are attempting to resolve a domain ID conflict, you must manually assign domain IDs. A disruptive restart is required to apply most configuration changes, including manually assigned domain IDs. Nondisruptive domain restarts are acceptable only when changing a preferred domain ID into a static one (and the actual domain ID remains the same).



Note

A static domain is specifically configured by the user and may be different from the runtime domain. If the domain IDs are different, the runtime domain ID changes to take on the static domain ID after the next restart, either disruptive or nondisruptive.



Tip

If a VSAN is in interop mode, you cannot disruptively restart the fcdomain for that VSAN.

You can apply most of the configurations to their corresponding runtime values. Each of the following sections provide further details on how the fcdomain parameters are applied to the runtime values.

Restarting a Domain

To restart the fabric disruptively or nondisruptively using Fabric Manager, perform this task:

- Step 1** Expand **Fabricxx > VSANxx**, and then choose **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to restart.

You see the Running tab configuration of the domain in the Information pane as shown in [Figure 11-2](#).

Figure 11-2 Running Domain Configuration

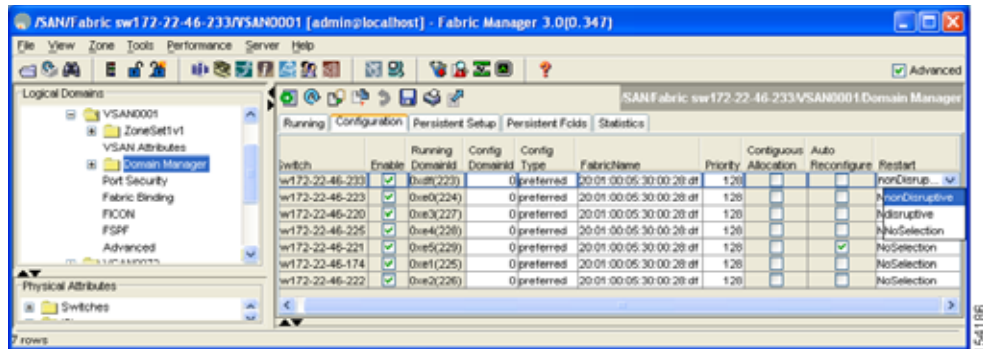
Switch	VSAN ID	State	DomainID	Local WWN	Local Priority	Principal WWN	Principal Priority
sw172-22-46-225	1	stable	Dxx4(228)	20:01:00:05:30:00:f1:e3	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-224	1	stable	Dxx6(230)	20:01:00:05:30:00:cb:57	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-223	1	stable	Dxx0(224)	20:01:00:05:30:00:61:df	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-222	1	stable	Dxx2(226)	20:01:00:05:30:00:eb:47	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-220	1	stable	Dxx3(227)	20:01:00:05:30:00:34:9f	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-174	1	stable	Dxx1(225)	20:01:00:05:30:01:9b:43	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-182	1	stable	Dxxa(234)	20:01:00:0d:ec:0e:94:c1	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-221	1	stable	Dxx5(229)	20:01:00:05:30:00:9a:5f	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-233	1	stable	Dxxf(223)	20:01:00:0d:ec:08:66:c1	128	Cisco 10:00:00:0d:ec:19:cb:0e	2

- Step 2** Click the **Configuration** tab.

You see the switch configuration as shown in [Figure 11-3](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 11-3 Configuring Domains



- Step 3** Choose **disruptive** or **nonDisruptive** in the Restart drop-down list for any switch in the fabric that you want to restart the fcdomain.
- Step 4** Click the **Apply Changes** icon to initiate the fcdomain restart.

About Switch Priority

By default, the configured priority is 128. The valid range to set the priority is between 1 and 254. Priority 1 has the highest priority. Value 255 is accepted from other switches, but cannot be locally configured.

Any new switch cannot become the principal switch when it joins a stable fabric. During the principal switch selection phase, the switch with the highest priority becomes the principal switch. If two switches have the same configured priority, the switch with the lower world-wide name (WWN) becomes the principal switch.

The priority configuration is applied to runtime when the fcdomain is restarted (see the [“About Domain Restart”](#) section on page 11-3). This configuration is applicable to both disruptive and nondisruptive restarts.

Configuring Switch Priority

To configure the priority for the principal switch using Fabric Manager, perform this task:

- Step 1** Expand **Fabricxx > VSANxx**, and then choose **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to set the principal switch priority for.

You see the domain’s running configuration in the Information pane as shown in [Figure 11-4](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 11-4 Running Domain Configuration

Switch	VSAN Id	State	Domain Id	Local WWN	Local Priority	Principal WWN	Principal Priority
sw172-22-46-225	1	stable	0xe4(228)	20:01:00:05:30:00:f1:e3	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-224	1	stable	0xe5(230)	20:01:00:05:30:00:cb:57	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-223	1	stable	0xe0(224)	20:01:00:05:30:00:61:df	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-222	1	stable	0xe2(226)	20:01:00:05:30:00:eb:47	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-220	1	stable	0xe3(227)	20:01:00:05:30:00:34:9f	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-174	1	stable	0xe1(225)	20:01:00:05:30:01:9b:43	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-182	1	stable	0xea(234)	20:01:00:0d:ec:0e:94:c1	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-221	1	stable	0xe5(229)	20:01:00:05:30:00:9a:5f	128	Cisco 10:00:00:0d:ec:19:cb:0e	2
sw172-22-46-233	1	stable	0xff(223)	20:01:00:0d:ec:08:66:c1	128	Cisco 10:00:00:0d:ec:19:cb:0e	2

Step 2 Click the **Configuration** tab.

You see the switch configuration as shown in [Figure 11-3](#).

Step 3 Set Priority to a high value for the switch in the fabric that you want to be the principal switch.

Step 4 Click the **Apply Changes** icon to save these changes.

About fcdomain Initiation

By default, the fcdomain feature is enabled on each switch. If you disable the fcdomain feature in a switch, that switch can no longer participate with other switches in the fabric. The fcdomain configuration is applied to runtime through a disruptive restart.

Enabling or Disabling fcdomains

To disable fcdomains in a single VSAN or a range of VSANs using Fabric Manager, perform this task:

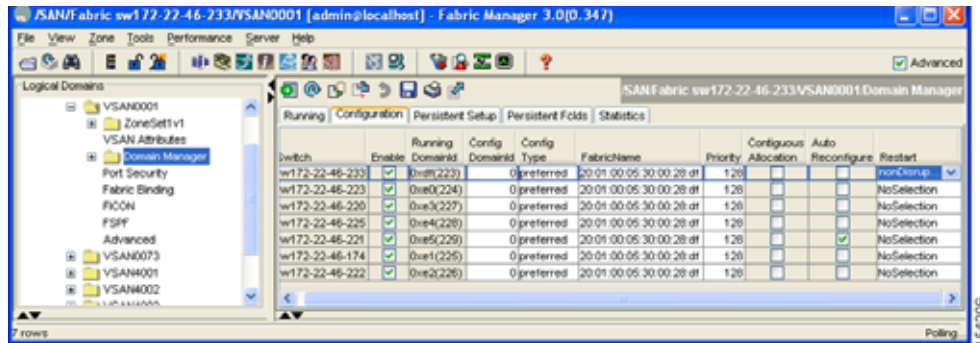
Step 1 Expand **Fabric:xx > VSAN:xx**, and then choose **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to disable fcdomain for.

You see the domain's running configuration in the Information pane.

Step 2 Click the **Configuration** tab and uncheck the **Enable** check box (see [Figure 11-5](#)) for each switch in the fabric that you want to disable fcdomain on.

Send comments to nx5000-docfeedback@cisco.com

Figure 11-5 Configuring Domains



Step 3 Click the **Apply Changes** icon to save these changes.

Setting Fabric Names

To set the fabric name value for a disabled fcdomain using Fabric Manager, perform this task:

- Step 1 Expand **Fabricxx > VSANxx**, and then choose **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to set the fabric name for.
You see the running configuration of the domain in the Information pane.
- Step 2 Click the **Configuration** tab and set the fabric name for each switch in the fabric.
- Step 3 Click the **Apply Changes** icon to save these changes.

About Incoming RCFs

You can choose to reject RCF request frames on a per-interface, per-VSAN basis. By default, the RCF reject option is disabled (that is, RCF request frames are not automatically rejected).

The RCF reject option takes effect immediately.

No fcdomain restart (see the [“About Domain Restart”](#) section on page 11-3).



Note

You do not need to configure the RFC reject option on virtual Fibre Channel interfaces, because these interfaces operate only in F port mode.

Rejecting Incoming RCFs

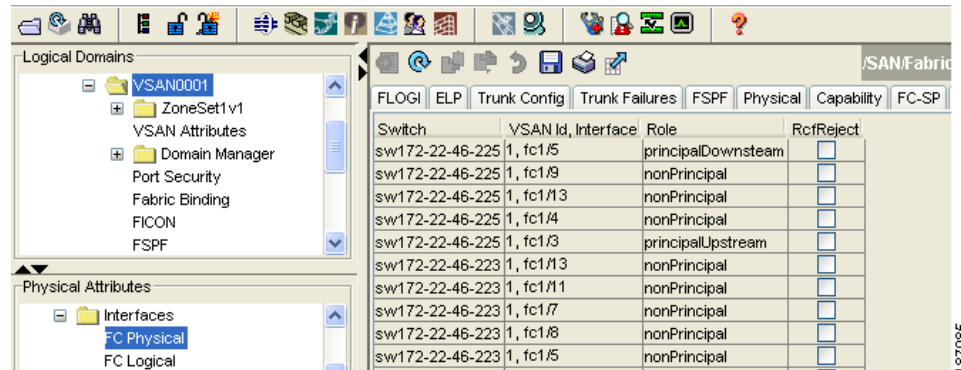
To reject incoming RCF request frames using Fabric Manager, perform this task:

- Step 1 In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**.
You see the Fibre Channel configuration in the Information pane.

Send comments to nx5000-docfeedback@cisco.com

- Step 2** Click the **Domain Mgr** tab.
You see the information in [Figure 11-6](#).

Figure 11-6 Rejecting Incoming RCF Request Frames



- Step 3** Check the **RcfReject** check box for each interface that you want to reject RCF request frames on.
Step 4 Click the **Apply Changes** icon to save these changes.

About Autoreconfiguring Merged Fabrics

By default, the autoreconfigure option is disabled. When you join two switches belonging to two different stable fabrics that have overlapping domains, the following situations can occur:

- If the autoreconfigure option is enabled on both switches, a disruptive reconfiguration phase is started.
- If the autoreconfigure option is disabled on either or both switches, the links between the two switches become isolated.

The autoreconfigure option takes immediate effect at runtime. You do not need to restart the fcdomain. If a domain is currently isolated due to domain overlap, and you later enable the autoreconfigure option on both switches, the fabric continues to be isolated. If you enabled the autoreconfigure option on both switches before connecting the fabric, a disruptive reconfiguration (RCF) will occur. A disruptive reconfiguration may affect data traffic. You can nondisruptively reconfigure the fcdomain by changing the configured domains on the overlapping links and eliminating the domain overlap.

Enabling Autoreconfiguration

To enable automatic reconfiguration in a specific VSAN (or range of VSANs) using Fabric Manager, perform this task:

- Step 1** Expand **Fabricxx > VSANxx**, and then choose **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to enable automatic reconfiguration for.
You see the running configuration of the domain in the Information pane.

Send comments to nx5000-docfeedback@cisco.com

- Step 2** Click the **Configuration** tab and check the **Auto Reconfigure** check box for each switch in the fabric that you want to automatically reconfigure.
- Step 3** Click the **Apply Changes** icon to save these changes.
-

Domain IDs

Domain IDs uniquely identify a switch in a VSAN. A switch may have different domain IDs in different VSANs. The domain ID is part of the overall FC ID.

This section describes how to configure domain IDs and includes the following topics:

- [About Domain IDs, page 11-8](#)
- [Specifying Static or Preferred Domain IDs, page 11-10](#)
- [About Allowed Domain ID Lists, page 11-10](#)
- [Configuring Allowed Domain ID Lists, page 11-11](#)
- [About CFS Distribution of Allowed Domain ID Lists, page 11-12](#)
- [Enabling Distribution, page 11-12](#)
- [Locking the Fabric, page 11-12](#)
- [Committing Changes, page 11-12](#)
- [Discarding Changes, page 11-13](#)
- [Clearing a Fabric Lock, page 11-13](#)
- [Displaying Pending Changes, page 11-14](#)
- [Displaying Session Status, page 11-14](#)
- [About Contiguous Domain ID Assignments, page 11-14](#)
- [Enabling Contiguous Domain ID Assignments, page 11-15](#)

About Domain IDs

The configured domain ID can be preferred or static. By default, the configured domain ID is 0 (zero) and the configured type is preferred.



Note

The 0 (zero) value can be configured only if you use the preferred option.

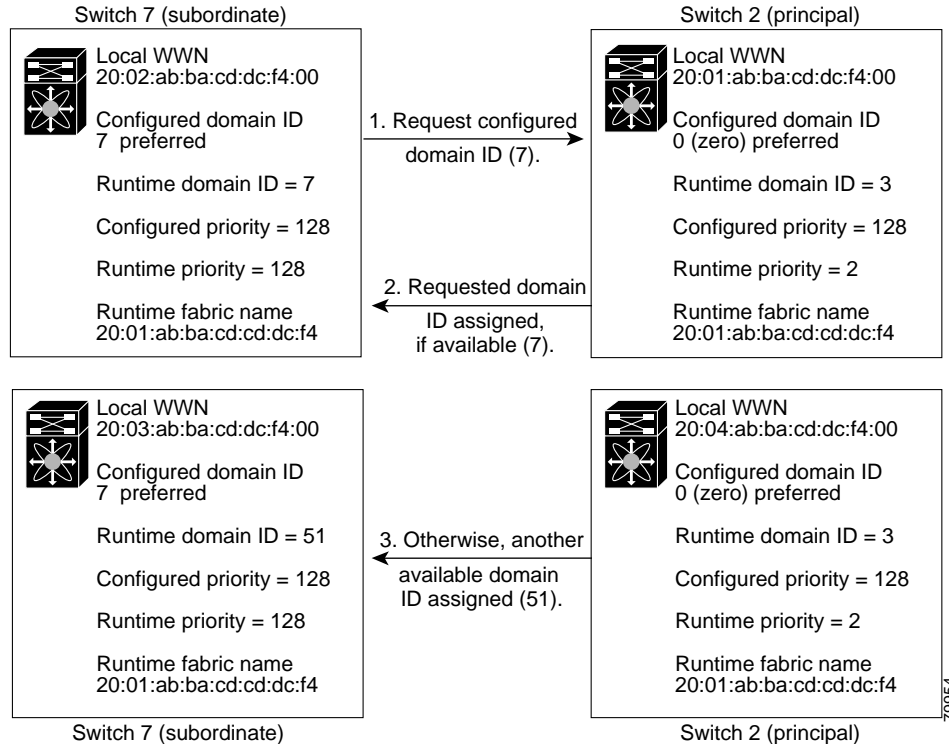
If you do not configure a domain ID, the local switch sends a random ID in its request. We recommend that you use static domain IDs.

When a subordinate switch requests a domain, the following process takes place (see [Figure 11-7](#)):

1. The local switch sends a configured domain ID request to the principal switch.
2. The principal switch assigns the requested domain ID if available. Otherwise, it assigns another available domain ID.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Figure 11-7 Configuration Process Using the Preferred Option



The operation of a subordinate switch changes based on three factors:

- The allowed domain ID lists.
- The configured domain ID.
- The domain ID that the principal switch has assigned to the requesting switch.

In specific situations, the changes are as follows:

- When the received domain ID is not within the allowed list, the requested domain ID becomes the runtime domain ID and all interfaces on that VSAN are isolated.
- When the assigned and requested domain IDs are the same, the preferred and static options are not relevant, and the assigned domain ID becomes the runtime domain ID.
- When the assigned and requested domain IDs are different, the following cases apply:
 - If the configured type is static, the assigned domain ID is discarded, all local interfaces are isolated, and the local switch assigns itself the configured domain ID, which becomes the runtime domain ID.
 - If the configured type is preferred, the local switch accepts the domain ID assigned by the principal switch and the assigned domain ID becomes the runtime domain ID.

If you change the configured domain ID, the change is only accepted if the new domain ID is included in all the allowed domain ID lists currently configured in the VSAN. Alternatively, you can also configure zero-preferred domain ID.



Caution

You must restart the **fcdomain** if you want to apply the configured domain changes to the runtime domain.

Send comments to nx5000-docfeedback@cisco.com

**Note**

If you have configured an allow domain ID list, the domain IDs that you add must be in that range for the VSAN. See the [“About Allowed Domain ID Lists”](#) section on page 11-10.

Specifying Static or Preferred Domain IDs

When you assign a static domain ID type, you are requesting a particular domain ID. If the switch does not obtain the requested address, it will isolate itself from the fabric. When you specify a preferred domain ID, you are also requesting a particular domain ID; however, if the requested domain ID is unavailable, then the switch will accept another domain ID.

While the static option can be applied at runtime after a disruptive or nondisruptive restart, the preferred option is applied at runtime only after a disruptive restart (see the [“About Domain Restart”](#) section on page 11-3).

**Note**

Within a VSAN all switches should have the same domain ID type (either static or preferred). If a configuration is mixed (some switches with static domain types and others with preferred), you may experience link isolation.

To specify a static or preferred domain ID using Fabric Manager, perform this task:

-
- Step 1** Expand **Fabricxx > VSANxx**, and then choose **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to configure the domain ID for.
- You see the running configuration of the domain in the Information pane.
- Step 2** Enter a value for the Config DomainID and click **static** or **preferred** from the Config Type drop-down list to set the domain ID for switches in the fabric.
- Step 3** Click the **Apply Changes** icon to save these changes.
-

About Allowed Domain ID Lists

By default, the valid range for an assigned domain ID list is from 1 to 239. You can specify a list of ranges to be in the allowed domain ID list and separate each range with a comma. The principal switch assigns domain IDs that are available in the locally configured allowed domain list.

Use allowed domain ID lists to design your VSANs with nonoverlapping domain IDs. This helps you in the future if you need to implement IVR without the NAT feature.

**Tip**

If you configure an allowed list on one switch in the fabric, we recommend that you configure the same list in all other switches in the fabric to ensure consistency or use CFS to distribute the configuration.

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.

Send comments to nx5000-docfeedback@cisco.com

- The locally configured domain ID of the switch must be in the allowed list.
- The intersection of the assigned domain IDs with other already configured domain ID lists must not be empty.

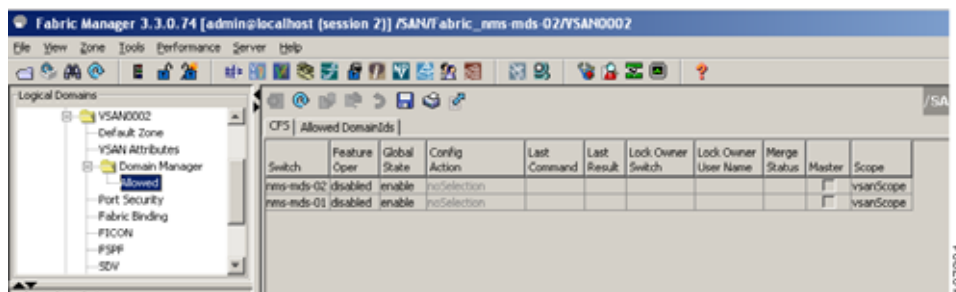
Configuring Allowed Domain ID Lists

To configure the allowed domain ID list using Fabric Manager, perform this task:

- Step 1** Expand **Fabricxx > VSANxx > Domain Manager**, and then choose **Allowed** in the Logical Domains pane for the fabric and VSAN for which you want to set the allowed domain ID list.

You see the CFS configuration in the Information pane (see [Figure 11-8](#)).

Figure 11-8 Allowed CFS Configuration Information



- Step 2** Set the Admin drop-down list to **enable** and set the Global drop-down list to **enable**.
- Step 3** Click **Apply Changes** to enable CFS distribution for the allowed domain ID list.
- Step 4** Click the **Allowed DomainIds** tab.

You see the Allowed Domain ID screen as shown in [Figure 11-9](#).

Figure 11-9 Allowed Domain ID List



- Step 5** Set the list to the allowed domain IDs list for this domain.
- Step 6** Click the **CFS** tab and set Config Action to **commit**.
- Step 7** Click the **Apply Changes** icon to commit this allowed domain ID list and distribute it throughout the VSAN.

Send comments to nx5000-docfeedback@cisco.com

About CFS Distribution of Allowed Domain ID Lists

You can enable the distribution of the allowed domain ID list configuration information to all Cisco SAN switches in the fabric using the Cisco Fabric Services (CFS) infrastructure. This feature allows you to synchronize the configuration across the fabric from the console of a single switch. Because the same configuration is distributed to the entire VSAN, you can avoid possible misconfiguration and the possibility that two switches in the same VSAN have configured incompatible allowed domains.

Use CFS to distribute the allowed domain ID list to ensure consistency in the allowed domain ID lists on all switches in the VSAN.



Note

We recommend configuring the allowed domain ID list and committing it on the principal switch.

For more information about CFS, see [Chapter 7, “Using Cisco Fabric Services.”](#)

Enabling Distribution

CFS distribution of allowed domain ID lists is disabled by default. You must enable distribution on all switches to which you want to distribute the allowed domain ID lists.

To enable (or disable) allowed domain ID list configuration distribution using Fabric Manager, perform this task:

-
- Step 1** Expand **Fabricxx > VSANxx > Domain Manager**, and then choose **Allowed** in the Logical Domains pane for the fabric and VSAN that you want to set the allowed domain ID list for.
- You see the CFS configuration in the Information pane.
- Step 2** Set the Admin drop-down list to **enable** and the Global drop-down list to **enable** to enable CFS distribution for the allowed domain ID list.
- Step 3** Click the **Apply Changes** icon to enable CFS distribution for the allowed domain ID list.
-

Locking the Fabric

The first action that modifies the existing configuration creates the pending configuration and locks the feature in the fabric. After you lock the fabric, the following conditions apply:

- No other user can make any configuration changes to this feature.
- A pending configuration is created by copying the active configuration. Subsequent modifications are made to the pending configuration and remain there until you commit the changes to the active configuration (and other switches in the fabric) or discard them.

Committing Changes

To apply the pending domain configuration changes to other SAN switches in the VSAN, you must commit the changes. The pending configuration changes are distributed and, on a successful commit, the configuration changes are applied to the active configuration in the SAN switches throughout the VSAN and the fabric lock is released.

Send comments to nx5000-docfeedback@cisco.com

To commit pending domain configuration changes and release the lock using Fabric Manager, perform this task:

-
- Step 1** Expand **Fabricxx > VSANxx > Domain Manager**, and then choose **Allowed** in the Logical Domains pane for the fabric and VSAN that you want to set the allowed domain ID list for.
- You see the CFS configuration in the Information pane.
- Step 2** Set the Config Action drop-down list to **commit**.
- Step 3** Click the **Apply Changes** icon to commit the allowed domain ID list and distribute it throughout the VSAN.
-

Discarding Changes

At any time, you can discard the pending changes to the domain configuration and release the fabric lock. If you discard (abort) the pending changes, the configuration remains unaffected and the lock is released.

To discard pending domain configuration changes and release the lock using Fabric Manager, perform this task:

-
- Step 1** Expand **Fabricxx > VSANxx > Domain Manager**, and then choose **Allowed** in the Logical Domains pane for the fabric and VSAN that you want to set the allowed domain ID list for.
- You see the CFS configuration in the Information pane.
- Step 2** Set the Config Action drop-down list to **abort**.
- Step 3** Click the **Apply Changes** icon to discard any pending changes to the allowed domain ID list.
-

Clearing a Fabric Lock

If you have performed a domain configuration task and have not released the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your pending changes are discarded and the fabric lock is released.



Tip

The pending changes are only available in the volatile directory and are discarded if the switch is restarted.

To release a fabric lock using Fabric Manager, perform this task:

-
- Step 1** Expand **Fabricxx > VSANxx > Domain Manager**, and then choose **Allowed** in the Logical Domains pane for the fabric and VSAN for which you want the allowed domain ID list.
- You see the CFS configuration in the Information pane.
- Step 2** Set the Config Action drop-down list to **clear**.

Send comments to nx5000-docfeedback@cisco.com

Step 3 Click the **Apply Changes** icon to clear the fabric lock.

Displaying Pending Changes

To display the pending configuration changes using Fabric Manager, perform this task:

Step 1 Expand **Fabricxx > VSANxx > Domain Manager > Allowed** in the Logical Domains pane for the fabric and VSAN that you want to set the allowed domain ID list for.

You see the CFS configuration in the Information pane.

Step 2 Set the Config View As drop-down list to **pending**.

Step 3 Click the **Apply Changes** icon to clear the fabric lock.

Step 4 Click the **AllowedDomainIds** tab.

You see the pending configuration for the allowed domain IDs list.

Displaying Session Status

To display the status of the distribution session using Fabric Manager, perform this task:

Step 1 Expand **Fabricxx > VSANxx > Domain Manager**, and then choose **Allowed** in the Logical Domains pane for the fabric and VSAN for which you want to set the allowed domain ID list.

Step 2 View the CFS configuration and session status in the Information pane.

About Contiguous Domain ID Assignments

By default, the contiguous domain assignment is disabled. When a subordinate switch requests the principal switch for two or more domains and the domains are not contiguous, the following situations can occur:

- If the contiguous domain assignment is enabled in the principal switch, the principal switch locates contiguous domains and assigns them to the subordinate switches. If contiguous domains are not available, the switch software rejects this request.
- If the contiguous domain assignment is disabled in the principal switch, the principal switch assigns the available domains to the subordinate switch.

Send comments to nx5000-docfeedback@cisco.com

Enabling Contiguous Domain ID Assignments

To enable contiguous domains in a specific VSAN (or a range of VSANs) using Fabric Manager, perform this task:

-
- Step 1** Expand **Fabricxx > VSANxx**, and then choose **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to enable contiguous domains for.
- You see the running configuration of the domain in the Information pane.
- Step 2** Click the **Configuration** tab and check the **Contiguous Allocation** check box for each switch in the fabric that will have contiguous allocation.
- Step 3** Click the **Apply Changes** icon to save these changes.
-

FC IDs

When an N port logs into a Cisco Nexus 5000 Series switch, it is assigned an FC ID. By default, the persistent FC ID feature is enabled. If this feature is disabled, the following situations can occur:

- An N port logs into a Cisco Nexus 5000 Series switch. The WWN of the requesting N 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.
- 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).

This section describes configuring FC IDs and includes the following topics:

- [About Persistent FC IDs, page 11-16](#)
- [Enabling the Persistent FC ID Feature, page 11-16](#)
- [Persistent FC ID Configuration Guidelines, page 11-16](#)
- [Configuring Persistent FC IDs, page 11-17](#)
- [About Unique Area FC IDs for HBAs, page 11-17](#)
- [Configuring Unique Area FC IDs for an HBA, page 11-18](#)
- [About Persistent FC ID Selective Purging, page 11-19](#)
- [Purging Persistent FC IDs, page 11-19](#)

Send comments to nx5000-docfeedback@cisco.com

About Persistent FC IDs

When persistent FC IDs are enabled, the following occurs:

- The current FC IDs in use 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

When persistent FC IDs are enabled, FC IDs cannot be changed after a reboot. FC IDs are enabled by default, but can be disabled for each VSAN.

A persistent FC ID assigned to an F port can be moved across interfaces and can continue to maintain the same persistent FC ID.

Enabling the Persistent FC ID Feature

To enable the persistent FC ID feature using Fabric Manager, perform this task:

-
- Step 1** Expand **Fabricxx > VSANxx**, and then choose **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to enable the Persistent FC ID feature for.
- You see the running configuration of the domain in the Information pane.
- Step 2** Click the **Persistent Setup** tab and check the **enable** check box for each switch in the fabric that will have persistent FC ID enabled.
- Step 3** Click the **Apply Changes** icon to save these changes.
-

Persistent FC ID Configuration Guidelines

When the persistent FC ID feature is enabled, you can enter the persistent FC ID submode and add static or dynamic entries in the FC ID database. By default, all added entries are static. Persistent FC IDs are configured on a per-VSAN basis.

When manually configuring a persistent FC ID, follow these requirements:

- 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.

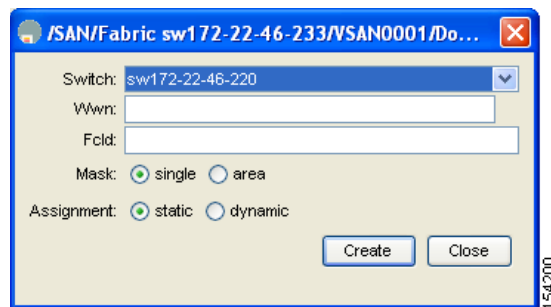
Send comments to nx5000-docfeedback@cisco.com

Configuring Persistent FC IDs

To configure persistent FC IDs using Fabric Manager, perform this task:

- Step 1** Expand **Fabricxx > VSANxx**, and then choose **Domain Manager** in the Logical Domains pane for the fabric and VSAN that you want to configure the Persistent FC ID list for.
You see the running configuration of the domain in the Information pane.
- Step 2** Click the **Persistent FcIds** tab and click **Create Row**.
You see the Create Persistent FC IDs dialog box as shown in [Figure 11-10](#).

Figure 11-10 Create Persistent FC IDs Dialog Box



- Step 3** Choose the switch, WWN, and FC ID that you want to make persistent.
- Step 4** Click either the **single** or **area** radio button in the Mask field.
- Step 5** Click either the **static** or **dynamic** radio button in the Assignment field.
- Step 6** Click the **Apply Changes** icon to save these changes.

About Unique Area FC IDs for HBAs



Note

Only read this section if the Host Bus Adapter (HBA) port and the storage port are connected to the same switch.

Some HBA ports require a different area ID than for the storage ports when they are both connected to the same switch. For example, if the storage port FC ID is 0x6f7704, the area for this port is 77. In this case, the HBA port's area can be anything other than 77. The HBA port's FC ID must be manually configured to be different from the storage port's FC ID.

Cisco Nexus 5000 Series switches 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.

Send comments to nx5000-docfeedback@cisco.com

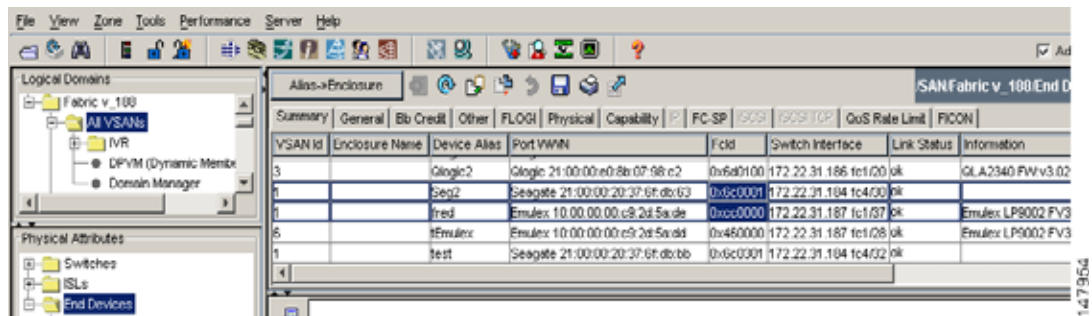
Configuring Unique Area FC IDs for an HBA

The following task uses an example configuration with a switch domain of 111(6f hex). The server connects to the switch over FCoE. The HBA port connects to interface vfc20/1 and the storage port connects to interface fc2/3 on the same switch.

To configure a different area ID for the HBA port using Fabric Manager, perform this task:

- Step 1** Expand **End Device** in the Physical Attributes pane and then choose the **Summary** tab to obtain the port WWN (Port Name field) of the HBA (see [Figure 11-11](#)).

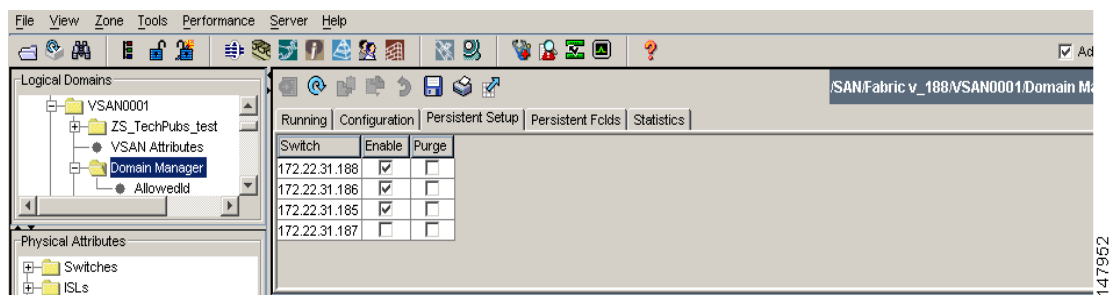
Figure 11-11 FLOGI Database Information in Fabric Manager



Note Both FC IDs in this setup have the same area 00 assignment.

- Step 2** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**. You see the Fibre Channel Interfaces information pane.
- Step 3** Set the Status Admin drop-down list to **down** for the interface that the HBA is connected to. This shuts down the HBA interface in the MDS switch.
- Step 4** Expand **Fabricxx > VSANxx**, and then choose **Domain Manager**.
- Step 5** Click the **Persistent Setup** tab in the Information pane to verify that the FC ID feature is enabled (see [Figure 11-12](#)).

Figure 11-12 Persistent FC ID Information in Fabric Manager



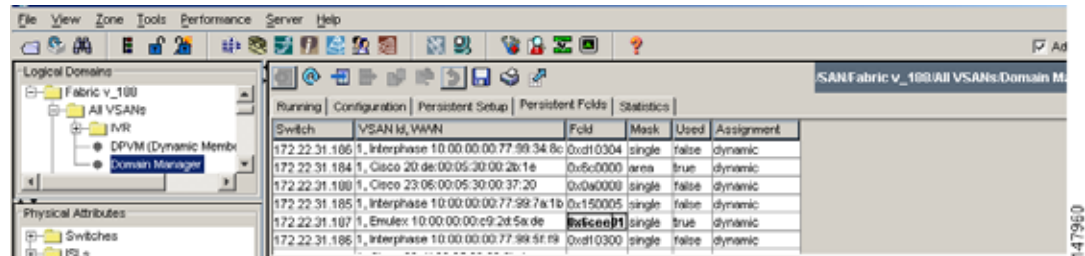
If this feature is disabled, continue with this procedure to enable persistent FC ID.

If this feature is already enabled, skip to [Step 7](#).

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

- Step 6** Check the **Enable** check box to enable the persistent FC ID feature in the switch (see [Figure 11-13](#)).
- Step 7** Click the **Persistent FCIDs** tab and assign a new FC ID with a different area allocation in the FcId field. For example, in [Figure 11-13](#) we replace *00* with *ee*.

Figure 11-13 Setting the FC ID in Fabric Manager



- Step 8** Click **Apply Changes** to save the new FC ID.
- Step 9** Compare the FC ID values to verify the FC ID of the HBA.



Note Both FC IDs now have different area assignments.

- Step 10** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**. Set the Status Admin drop-down list to **up** for the interface that the HBA is connected to. This enables the HBA interface in the Cisco Nexus 5000 Series switch.

About Persistent FC ID Selective Purging

Persistent FC IDs can be purged selectively. Static entries and FC IDs currently in use cannot be deleted. [Table 11-1](#) identifies the FC ID entries that are deleted or retained when persistent FC IDs are purged.

Table 11-1 Purged FC IDs

Persistent FC ID state	Persistent Usage State	Action
Static	In use	Not deleted
Static	Not in use	Not deleted
Dynamic	In use	Not deleted
Dynamic	Not in use	Deleted

Purging Persistent FC IDs

To purge persistent FC IDs using Fabric Manager, perform this task:

Send comments to nx5000-docfeedback@cisco.com

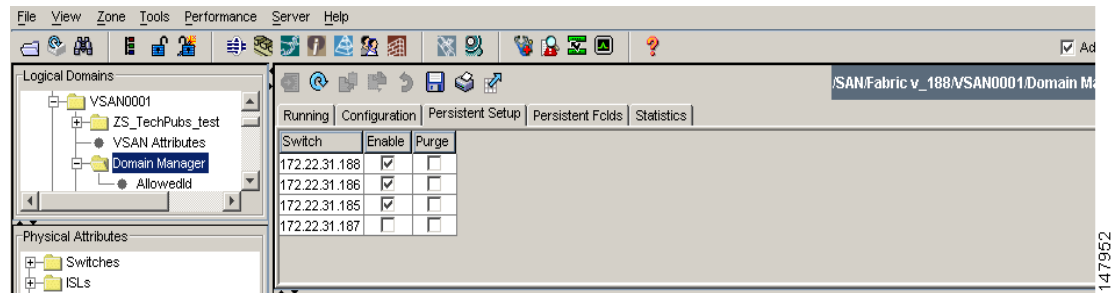
Step 1 Expand **Fabricxx > All VSANs > Domain Manager** in the Logical Domains pane for the fabric that you want to purge the Persistent FC IDs for.

You see the running configuration of the domain in the Information pane.

Step 2 Click the **Persistent Setup** tab.

You see the persistent FC ID setup in the Information pane as shown in [Figure 11-14](#).

Figure 11-14 Persistent FC ID Information in Fabric Manager



Step 3 Check the **Purge** check box for the switch that you want to purge persistent FC IDs on (see [Figure 11-14](#)).

Step 4 Click the **Apply Changes** icon to save these changes.

Displaying fcdomain Statistics

Fabric Manager collects statistics for fcdomain and displays them in the Information pane.

To display fcdomain statistics using Fabric Manager, perform this task:

Step 1 Expand **Fabricxx > All VSANs**, and then choose **Domain Manager** in the Logical Domains pane for the fabric that you want to display statistics for.

You see the running configuration of the domain in the Information pane.

Step 2 Click the **Statistics** tab.

You see the FC ID statistics in the Information pane.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Default Settings

Table 11-2 lists the default settings for all fcdomain parameters.

Table 11-2 *Default fcdomain Parameters*

Parameters	Default
fcdomain feature	Enabled
Configured domain ID	0 (zero)
Configured domain	Preferred
autoreconfigure option	Disabled
contiguous-allocation option	Disabled
Priority	128
Allowed list	1 to 239
Fabric name	20:01:00:05:30:00:28:df
rcf-reject	Disabled
Persistent FC ID	Enabled
Allowed domain ID list configuration distribution	Disabled

Send comments to nx5000-docfeedback@cisco.com



Configuring N-Port Virtualization

N-port virtualization (NPV) reduces the number of Fibre Channel domain IDs used in a SAN fabric. Edge switches operating in NPV mode do not join a fabric; they pass traffic between the NPV core switch and the end devices, which eliminates the need for a unique domain ID in each edge switch.

This chapter includes the following sections:

- [Information About NPV, page 12-1](#)
- [Guidelines and Limitations, page 12-4](#)
- [Configuring NPV, page 12-4](#)

Information About NPV

Typically, Fibre Channel networks are deployed using a core-edge model with a large number of fabric switches connected to core devices. However, as the number of ports in the fabric increases, the number of switches deployed also increases, resulting in a dramatic increase in the number of domain IDs (the maximum number supported in one SAN is 239). This challenge becomes even more difficult when a large number of blade switches are deployed in a Fibre Channel network.

NPV solves the increase in the number of domain IDs by sharing the domain ID of the NPV core switch among multiple NPV switches.

The NPV edge switch aggregates multiple locally connected N ports into one or more external NP links. The edge switch appears as a host to the core Fibre Channel switch, and as a Fibre Channel switch to the servers in the fabric switch or blade switch.


NPV reduces the need for additional ports on the core switch because multiple devices attach to the same port on the NPV core switch.

[Figure 12-1](#) shows an interface-level view of an NPV configuration.

In Cisco Nexus 5000 Series switches, physical Fibre Channel interfaces can be NP ports or F ports. Virtual Fibre Channel interfaces can be F ports.

Send comments to nx5000-docfeedback@cisco.com

Figure 12-1 Cisco NPV Configuration–Interface View



Note

In-order data delivery is not required in NPV mode because the exchange between two end devices always takes the same uplink to the core from the NPV device. For traffic beyond the NPV device, core switches will enforce in-order delivery if configured.

Switch operation in NPV mode is described in the following topics:

- [NP Ports, page 12-2](#)
- [NP Links, page 12-2](#)
- [FLOGI Operation, page 12-2](#)

NP Ports

An NP port (proxy N port) is a port on a switch that is in NPV mode and connected to the core NPV switch through an F port. NP ports operate as N ports that function as proxies for multiple physical N ports.

NP Links

An NP link is basically an NPIV uplink to a specific end device. NP links are established when the uplink to the NPV core switch comes up; the links are terminated when the uplink goes down. Once the uplink is established, the NPV switch performs an internal FLOGI to the NPV core switch, and then (if the FLOGI is successful) registers itself with the NPV core switch's name server. Subsequent FLOGIs from end switches in this NP link are converted to FDISCs.

Server links are uniformly distributed across the NP links. All the end devices behind a server link will be mapped to only one NP link.

FLOGI Operation

When an NP port comes up, the Cisco Nexus 5000 Series switch first logs itself in to the NPV core switch and sends a FLOGI request that includes the following parameters:

- The fabric port WWN (fWWN) of the NP port used as the pWWN in the internal login.
- The VSAN-based switch WWN (sWWN) of the Cisco Nexus 5000 Series switch used as node WWN (nWWN) in the internal FLOGI.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

After completing its FLOGI request, the Cisco Nexus 5000 Series switch registers itself with the fabric name server using the following additional parameters:

- Switch name and interface name (for example, fc2/4) of the NP port is embedded in the symbolic port name in the name server registration of the NPV device itself.
- The IP address of the Cisco Nexus 5000 Series switch is registered as the IP address in the name server registration of the NPV device.



Note

The BB_SCN of internal FLOGIs on NP ports is always set to zero. The BB_SCN is supported at the F-port of the NPV device.

Figure 12-2 shows the internal FLOGI flows between an NPV core switch and an NPV device.

Figure 12-2 Internal FLOGI Flows




Table 12-1 identifies the internal FLOGI parameters that appear in Figure 12-2.

Table 12-1 Internal FLOGI Parameters

Parameter	Derived From
pWWN	The fWWN of the NP port.
nWWN	The VSAN-based sWWN of the NPV device.
fWWN	The fWWN of the F port on the NPV core switch.
symbolic port name	The switch name and NP port interface string. Note If there is no switch name available, then the output will read “switch.” For example, switch: fc2/3.
IP address	The IP address of the NPV device.
symbolic node name	The NPV switch name.

Although fWWN-based zoning is supported for NPV devices, it is not recommended because of these factors:

- Zoning is not enforced at the NPV device (rather, it is enforced on the NPV core switch).
- Multiple devices attached to an NPV device log in through the same F port on the core, so they cannot be separated into different zones.

Send comments to nx5000-docfeedback@cisco.com

- The same device might log in using different fWWNs on the core switch (depending on the NPV link it uses) and may need to be zoned using different fWWNs.

Guidelines and Limitations

The following are recommended guidelines and requirements when deploying NPV:

- You can configure zoning for end devices that are connected to NPV devices using all available member types on the NPV core switch. If fWWN, sWWN, domain, or port-based zoning is used, then fWWN, sWWN, domain or port of the NPV core switch should be used.
- Port tracking is supported in NPV mode. See the [“Information About Port Tracking” section on page 27-1](#).
- Port security is supported on the NPV core switch for devices logged in through the NPV switch. Port security is enabled on the NPV core switch on a per-interface basis. To enable port security on the NPV core switch for devices logging in through an NPV switch, you must adhere to the following requirements:
 - The internal FLOGI must be in the port security database; in this way, the port on the NPV core switch will allow communications and links.
 - All the end device pWWNs must also be in the port security database.

By grouping devices into different NPV sessions based on VSANs, it is possible to support multiple VSANs at the NPV-enabled switch. The correct uplink must be selected based on the VSANs that the uplink can carry.

- NPV uses a load-balancing algorithm to automatically assign end devices in a VSAN to one of the NPV core switch links (in the same VSAN) upon initial login. If there are multiple NPV core switch links in the same VSAN, then you cannot assign an end device to a specific core switch link.
- If a server interface goes down and then returns to service, the interface may not be assigned to the same core switch link.
- The server interface is only operational when its assigned core switch link is operational.
- Both servers and targets can be connected to the switch when in NPV mode.
- Local switching is not supported; all traffic is switched in the NPV core switch.
- NPV devices can connect to multiple NPV core switches. In other words, different NP ports can be connected to different NPV core switches.
- NPV supports NPIV-capable module servers (nested NPIV).
- Only F, NP, and SD ports are supported in NPV mode.

Configuring NPV

When you enable NPV, your system configuration is erased and the system is rebooted with NPV mode enabled.



Note

We recommend that you save your current configuration either in boot flash memory or to a TFTP server before NPV (if the configuration is required for later use).

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Configuring NPV with Device Manager

To use Device Manager to configure NPV, perform this task:

-
- Step 1** Launch Device Manager from the Cisco Nexus 5000 Series switch to enable NPV.
 - Step 2** From the Admin drop-down menu, choose **Feature Control**. In the **Action** field, choose **enable** for the NPV feature and click **Apply**.
 - Step 3** From the Interface drop-down list, choose **FC All** to configure the external interfaces on the NPV device.
 - Step 4** In the Mode Admin column, choose the **NP** port mode for each external interface and click **Apply**.
 - Step 5** To configure the server interfaces on the Cisco Nexus 5000 Series switch, from the Interface drop-down list, choose **FC All**.
 - Step 6** In the Mode Admin column, choose **F** port mode for each server interface and click **Apply**.
 - Step 7** The default Admin status is **down**. After configuring port modes, you must choose **up** Admin Status to bring up the links.
-

Send comments to nx5000-docfeedback@cisco.com



Configuring VSAN Trunking

This chapter describes the VSAN trunking feature provided in Cisco Nexus 5000 Series switches.

This chapter includes the following sections:


- [Information About VSAN Trunking, page 13-1](#)
- [Configuring VSAN Trunking, page 13-3](#)
- [Default Settings, page 13-7](#)

Information About VSAN Trunking

VSAN trunking enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link, using enhanced ISL (EISL) frame format (see [Figure 13-1](#)).

VSAN trunking is supported on native Fibre Channel interfaces, but not on virtual Fibre Channel interfaces.

Figure 13-1 VSAN Trunking



The VSAN 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.

Additional information about VSAN trunking is covered in the following topics:


- [VSAN Trunking Mismatches, page 13-2](#)
- [VSAN Trunking Protocol, page 13-2](#)

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

VSAN Trunking Mismatches


If you misconfigure VSAN configurations across E ports, issues can occur such as the merging of traffic in two VSANs (causing both VSANs to mismatch). The VSAN trunking protocol validates the VSAN interfaces at both ends of an ISL to avoid merging VSANs (see [Figure 13-2](#)).

Figure 13-2 VSAN Mismatch



In this example, the trunking protocol detects potential VSAN merging and isolates the ports involved. The trunking protocol cannot detect merging of VSANs when a third-party switch is placed in between two Cisco Nexus 5000 Series switches (see [Figure 13-3](#)).

Figure 13-3 Third-Party Switch VSAN Mismatch



VSAN 2 and VSAN 3 are effectively merged with overlapping entries in the name server and the zone applications. The Cisco MDS 9000 Fabric Manager helps detect such topologies.

VSAN Trunking Protocol

The trunking protocol is important for E-port and TE-port operations. It supports the following capabilities:

- 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 VSAN trunking protocol is enabled. If the trunking protocol is disabled on a switch, no port on that switch can apply new trunk configurations. Existing trunk configurations are not affected: the TE port continues to function in trunk mode, but only supports traffic in VSANs that it negotiated with previously (when the trunking protocol was enabled). Other switches that are directly connected to this switch are similarly affected on the connected interfaces. If you need to merge traffic from different port VSANs across a nontrunking ISL, disable the trunking protocol.

Send comments to nx5000-docfeedback@cisco.com

Configuring VSAN Trunking

This section explains how to configure VSAN trunking and includes the following topics:

- [Guidelines and Restrictions, page 13-3](#)
- [About Trunk Mode, page 13-3](#)
- [Configuring Trunk Mode, page 13-4](#)
- [About Trunk-Allowed VSAN Lists, page 13-5](#)
- [Configuring an Allowed-Active List of VSANs, page 13-6](#)

Guidelines and Restrictions

When configuring VSAN trunking, note the following guidelines:

- We recommend that both ends of a VSAN trunking ISL belong to the same port VSAN. On platforms or fabric switches where the port VSANs are different, one end returns an error, and the other is not connected.
- To avoid inconsistent configurations, shut all E ports before enabling or disabling the VSAN trunking protocol.

About Trunk Mode

By default, trunk mode is enabled in all Fibre Channel interfaces. However, trunk mode configuration takes effect only in E-port mode. You can configure trunk mode as on (enabled), off (disabled), or auto (automatic). The default trunk mode is on. The trunk mode configurations at the two ends of the link determine the trunking state of the link and the port modes at both ends (see [Table 13-1](#)).

Table 13-1 Trunk Mode Status Between Switches

Your Trunk Mode Configuration		Resulting State and Port Mode	
Switch 1	Switch 2	Trunking State	Port Mode
On	Auto or on	Trunking (EISL)	TE port
Off	Auto, on, or off	No trunking (ISL)	E port
Auto	Auto	No trunking (ISL)	E port



Tip

The preferred configuration on the Cisco Nexus 5000 Series switches is that one side of the trunk is set to auto and the other is set to on.



Note

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

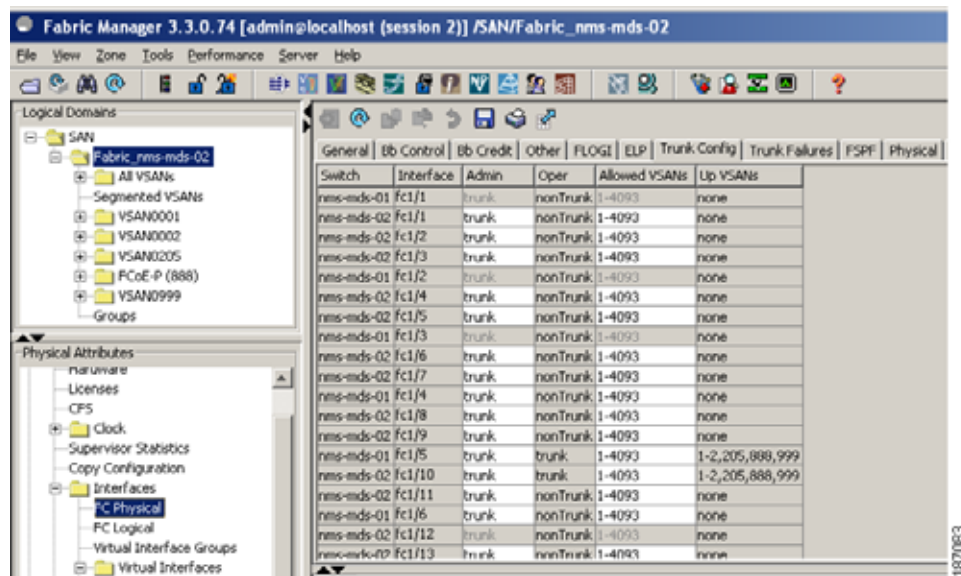
Send comments to nx5000-docfeedback@cisco.com

Configuring Trunk Mode

To configure trunk mode using Fabric Manager, perform this task:


- Step 1 In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**.
You see the interface configuration in the Information pane.
- Step 2 Click the **Trunk Config** tab to modify the trunking mode for the selected interface.
You see the information shown in [Figure 13-4](#).

Figure 13-4 Trunking Configuration



- Step 3 Make changes to the Admin and Allowed VSANs values.
- Step 4 Click the **Trunk Failures** tab to check the failure state of an ISL.
You see the reason listed in the FailureCause column (see [Figure 13-5](#)).

Figure 13-5 Trunk Failures Tab



- Step 5 Click the **Apply Changes** icon.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)


About Trunk-Allowed VSAN Lists

Each Fibre Channel interface has an associated trunk-allowed VSAN list. In TE-port mode, frames are transmitted and received in one or more VSANs specified in this list. By default, the complete 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 13-6](#), switch 1 has VSANs 1 through 5, switch 2 has VSANs 1 through 3, and switch 3 has VSANs 1, 2, 4, and 5 with a default configuration of trunk-allowed VSANs. All VSANs configured in all three switches are allowed-active. However, only the common set of allowed-active VSANs at the ends of the ISL become operational as shown in [Figure 13-6](#).

Figure 13-6 Default Allowed-Active VSAN Configuration



You can configure a selected set of VSANs (from the allowed-active list) to control access to the VSANs specified in a trunking ISL.


Using [Figure 13-6](#) as an example, you can configure the list of allowed VSANs on a per-interface basis (see [Figure 13-7](#)). For example, if VSANs 2 and 4 are removed from the allowed VSAN list of ISLs connecting to switch 1, the operational allowed list of VSANs for each ISL would be as follows:

- The ISL between switch 1 and switch 2 includes VSAN 1 and VSAN 3.
- The ISL between switch 2 and switch 3 includes VSAN 1 and VSAN 2.
- The ISL between switch 3 and switch 1 includes VSAN 1, 2, and 5.

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

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Figure 13-7 Operational and Allowed VSAN Configuration



Configuring an Allowed-Active List of VSANs

To configure an allowed-active list of VSANs for an interface using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Interfaces**, and then choose **FC Physical**.
You see the interface configuration in the Information pane.
 - Step 2** Click the **Trunk Config** tab.
You see the current trunk configuration.
 - Step 3** Set Allowed VSANs to the list of allowed VSANs for each interface that you want to configure.
 - Step 4** Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.
-

Send comments to nx5000-docfeedback@cisco.com

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

Send comments to nx5000-docfeedback@cisco.com



Configuring SAN Port Channels

SAN port channels refer to the aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy.

On Cisco Nexus 5000 Series switches, SAN port channels can include physical Fibre Channel interfaces, but not virtual Fibre Channel interfaces. A SAN port channel can include up to eight Fibre Channel interfaces.

This chapter discusses the SAN port channel feature provided in the switch and includes the following sections:

- [Information About SAN Port Channels, page 14-1](#)
- [Configuring SAN Port Channels, page 14-5](#)
- [Interfaces in a SAN Port Channel, page 14-12](#)
- [Port Channel Protocol, page 14-15](#)
- [Verifying SAN Port Channel Configuration, page 14-19](#)
- [Default Settings, page 14-19](#)

Information About SAN Port Channels

A SAN port channel has the following functionality:

- Provides a point-to-point connection over ISL (E ports) or EISL (TE ports). Multiple links can be combined into a SAN port channel.
- 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 SAN port channel, the upper layer protocol is not aware of it. To the upper layer protocol, the link is still there, although the bandwidth is diminished. The routing tables are not affected by link failure.

Cisco Nexus 5000 Series switches support a maximum of four SAN port channels (with eight interfaces per port channel). A port channel number refers to the unique (within each switch) identifier associated with each channel group. This number ranges from 1 to 256.

Send comments to nx5000-docfeedback@cisco.com

This section describes SANs and includes the following topics:

- [Understanding Port Channels and VSAN Trunking, page 14-2](#)
- [Understanding Load Balancing, page 14-3](#)


Understanding Port Channels and VSAN Trunking

Switches in the Cisco Nexus 5000 Series implement VSAN trunking and port channels as follows:

- A SAN port channel enables several physical links to be combined into one aggregated logical link.
- An industry standard E port can link to other vendor switches and is referred to as inter-switch link (ISL), as shown on the left side of [Figure 14-1](#).
- VSAN trunking enables a link transmitting frames in the EISL format to carry traffic for multiple VSAN. When trunking is operational on an E port, that E port becomes a TE port. EISLs connects only between Cisco switches, as shown on the right side of [Figure 14-1](#).

See [Chapter 13, “Configuring VSAN Trunking”](#) for information on trunk interfaces.


Figure 14-1 VSAN Trunking Only



You can create a SAN port channel with members that are E ports, as shown on the left side of [Figure 14-2](#). In this configuration, the port channel implements a logical ISL (carrying traffic for one VSAN).

You can create a SAN port channel with members that are TE-ports, as shown on the right side of [Figure 14-2](#). In this configuration, the port channel implements a logical EISL (carrying traffic for multiple VSANs).

Figure 14-2 Port Channels and VSAN Trunking



[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

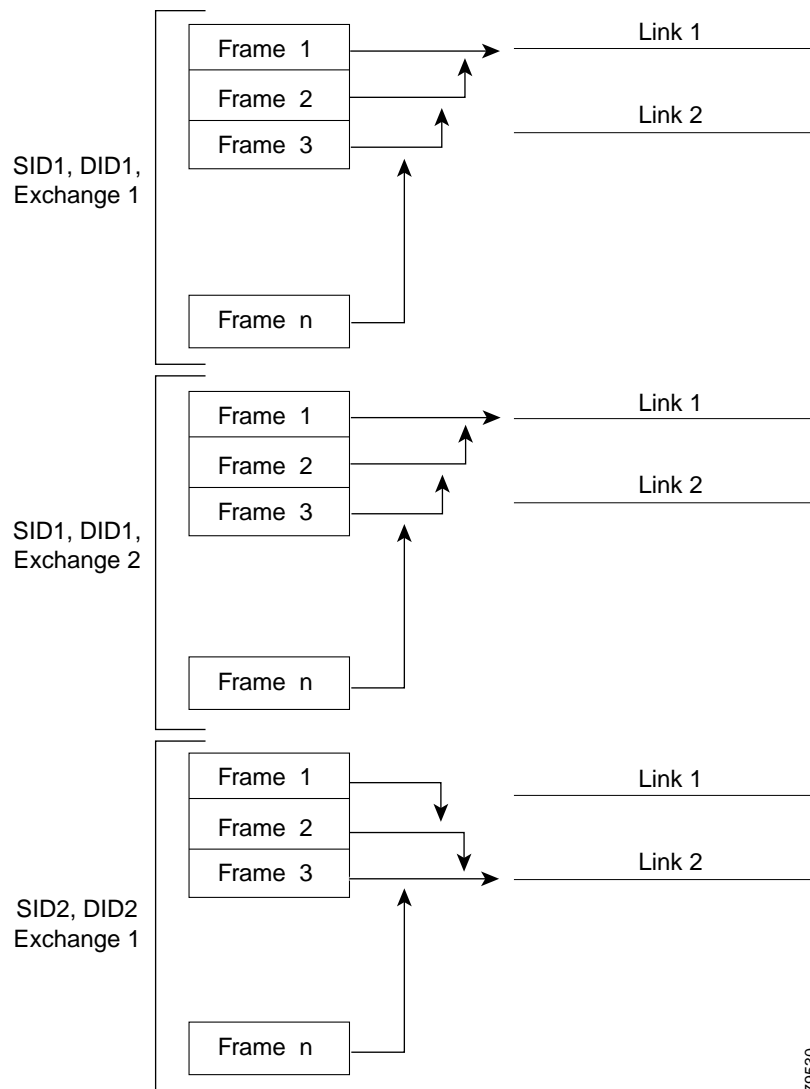
Understanding Load Balancing

Load-balancing functionality can be provided using the following methods:

- 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 is assigned to a link, and then subsequent frames in the exchange follow the same link. However, subsequent exchanges can use a different link. This method provides finer granularity for load balancing while preserving the order of frames for each exchange.

Figure 14-3 illustrates how flow-based 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.

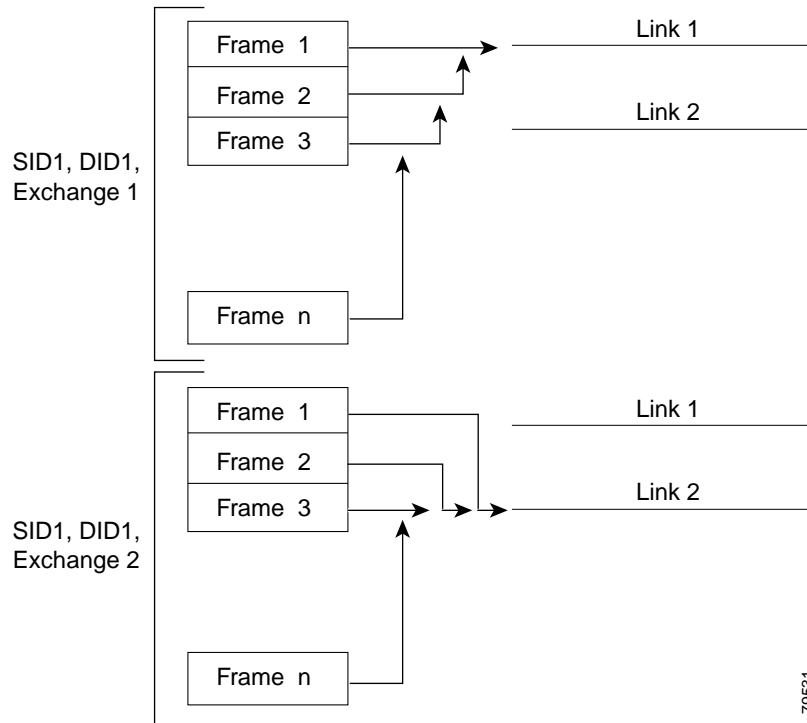
Figure 14-3 SID1, DID1, and Flow-Based Load Balancing



[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

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



[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Configuring SAN Port Channels

SAN port channels are created with default values. You can change the default configuration just as any other physical interface.

Figure 14-5 provides examples of valid SAN port channel configurations.

Figure 14-5 Valid SAN Port Channel Configurations

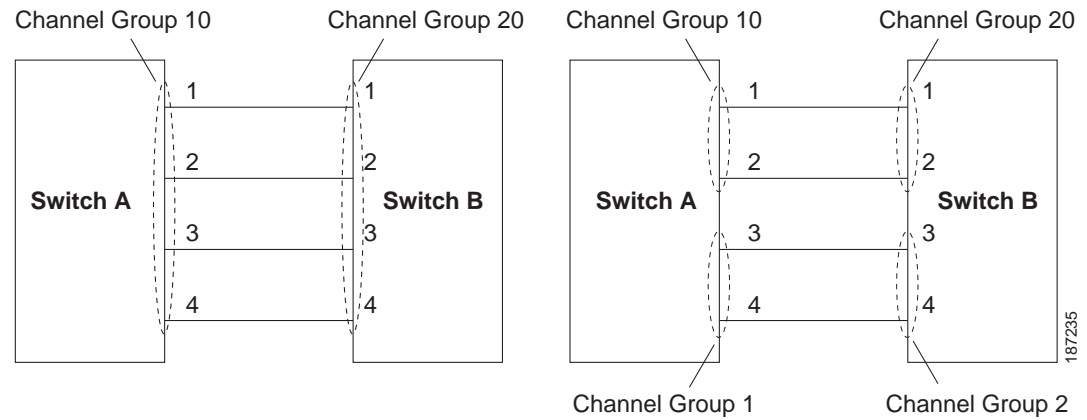
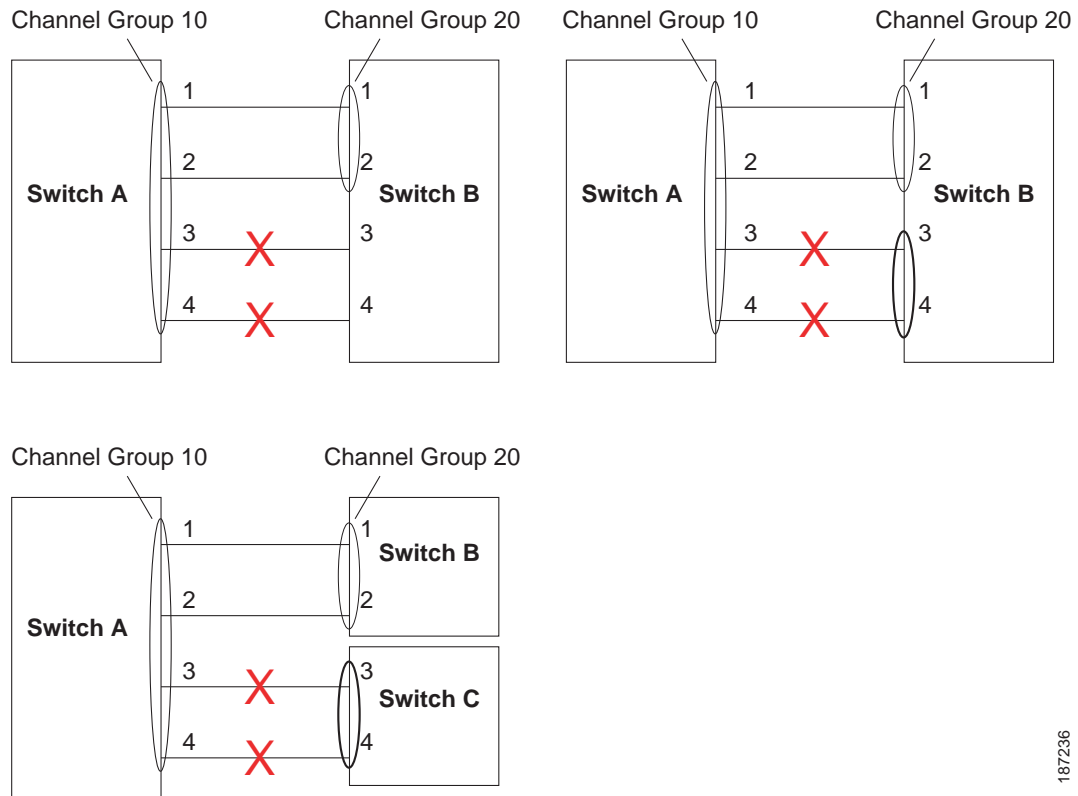


Figure 14-6 shows examples of invalid configurations. Assuming that the links are brought up in the 1, 2, 3, 4 sequence, links 3 and 4 will be operationally down as the fabric is misconfigured.

Send comments to nx5000-docfeedback@cisco.com

Figure 14-6 Misconfigured Configurations



187236

This section shows how to configure and modify SAN port channels and includes the following topics:

- [SAN Port Channel Configuration Guidelines, page 14-6](#)
- [Configuring SAN Port Channels, page 14-7](#)
- [About SAN Port Channel Modes, page 14-10](#)
- [About SAN Port Channel Deletion, page 14-11](#)
- [Deleting SAN Port Channels, page 14-11](#)

SAN Port Channel Configuration Guidelines

Before configuring a SAN port channel, consider the following guidelines:

- Configure the SAN port channel using Fibre Channel ports from both expansion modules to provide increased availability (if one of the expansion modules failed).
- Ensure that one SAN port channel is not connected to different sets of switches. SAN port channels require point-to-point connections between the same set of switches.

If you misconfigure SAN port channels, you may receive a misconfiguration message. If you receive this message, the port channel's physical links are disabled because an error has been detected.

Send comments to nx5000-docfeedback@cisco.com

If the following requirements are not met, a SAN port channel error is detected:

- Each switch on either side of a SAN port channel must be connected to the same number of interfaces.
- Each interface must be connected to a corresponding interface on the other side (see [Figure 14-6](#) for an example of an invalid configuration).
- Links in a SAN port channel cannot be changed after the port channel is configured. If you change the links after the port channel is configured, be sure to reconnect the links to interfaces within the port channel and reenables the links.


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

Configuring SAN Port Channels

To create a SAN port channel using the Port Channel Wizard in Fabric Manager, perform this task:

- Step 1** Click the **Port Channel Wizard** icon in the toolbar (see [Figure 14-7](#)).

Figure 14-7 Port Channel Wizard Icon



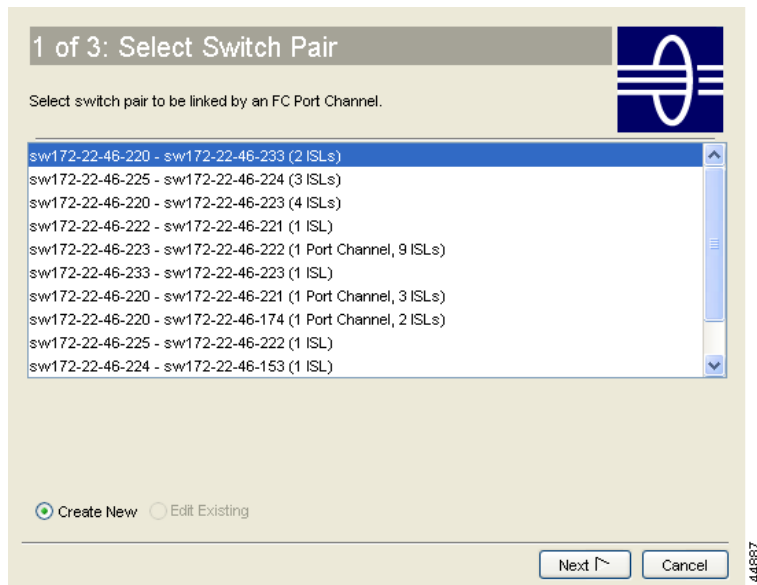
You see the first Port Channel Wizard screen.

- Step 2** Choose a switch pair.

[Figure 14-8](#) shows a list of the switch pairs.

Send comments to nx5000-docfeedback@cisco.com

Figure 14-8 Select Switch Pairs

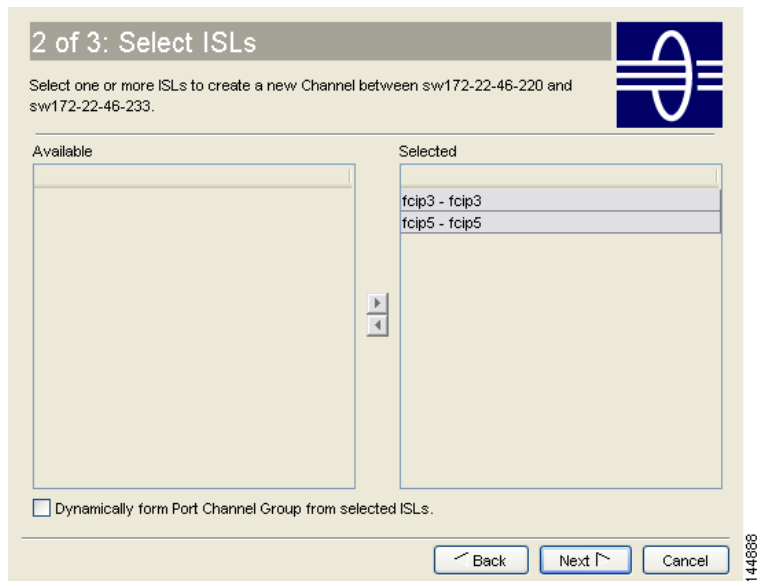


Step 3 Click **Next**.

Step 4 Select the ISLs.

Figure 14-9 shows a list of the ISLs.

Figure 14-9 Select ISLs



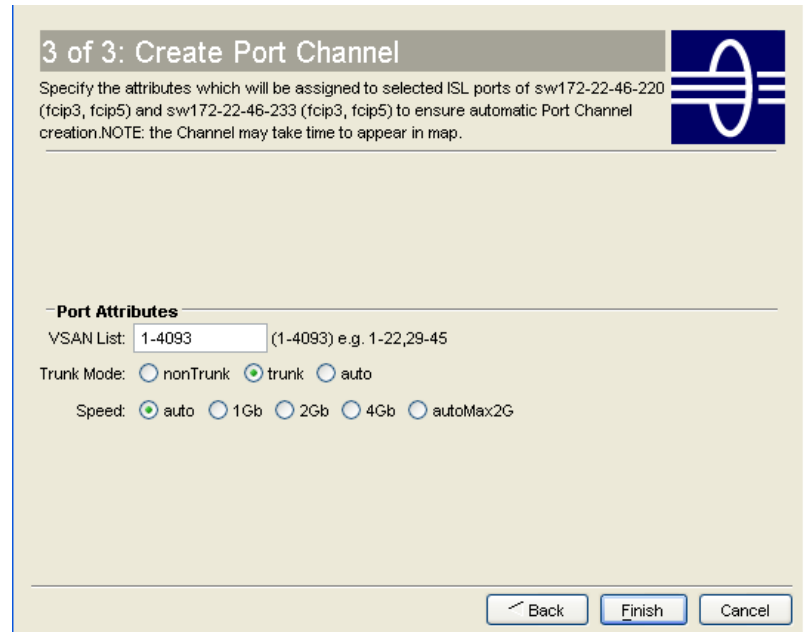
Step 5 (Optional) Check the **Dynamically form Port Channel Group from selected ISLs** check box if you want to dynamically create the SAN port channel and make the ISL properties identical for the Admin, Trunk, Speed, and VSAN attributes.

Step 6 Click **Next**.

Send comments to nx5000-docfeedback@cisco.com


- Step 7** If you chose to dynamically form a SAN port channel from selected ISLs, you see the final Port Channel Wizard screen (see [Figure 14-10](#)). Set the VSAN List, Trunk Mode, and Speed and proceed to [Step 11](#).

Figure 14-10 Dynamically Form a Port Channel



- Step 8** If you did not choose to dynamically form a SAN port channel, you see the third Port Channel Wizard dialog box (see [Figure 14-11](#)).

Figure 14-11 Create a Port Channel



- Step 9** Change the channel ID or description for each switch, if necessary.
- Step 10** Review the attributes at the bottom of the screen, and set them if applicable.

The following attributes are shown in [Figure 14-11](#):

- VSAN List—A list of VSANs to which the ISLs belong.
- Trunk Mode—You can enable trunking on the links in the SAN port channel. Choose **trunking** if your link is between TE ports. Choose **nontrunking** if your link is between E ports. Choose **auto** if you are not sure.
- Force Admin, Trunk, Speed, and VSAN attributes to be identical—This check box ensures that the same parameter settings are used in all physical ports in the channel. If these settings are not identical, the ports cannot become part of the SAN port channel.
- Speed—The port speed values are **auto**, **1Gb**, **2Gb**, **4Gb**, and **autoMax2G**.

- Step 11** Click **OK**.

Send comments to nx5000-docfeedback@cisco.com

The SAN port channel is created. It may take a few minutes before the new port channel is visible in the Fabric pane.

About SAN Port Channel Modes

You can configure each SAN port channel with a channel group mode parameter to determine the port channel protocol behavior for all member ports in this channel group. The possible values for a channel group mode are as follows:

- **On (default)**—The member ports only operate as part of a SAN port channel or remain inactive. In this mode, the port channel protocol is not initiated. However, if a port channel protocol frame is received from a peer port, the software indicates its nonnegotiable status. Port channels configured in the On mode require you to explicitly enable and disable the port channel member ports at either end if you add or remove ports from the port channel configuration. You must physically verify that the local and remote ports are connected to each other.
- **Active**—The member ports initiate port channel protocol negotiation with the peer port(s) regardless of the channel group mode of the peer port. If the peer port, while configured in a channel group, does not support the port channel protocol, or responds with a nonnegotiable status, it will default to the On mode behavior. The Active port channel mode allows automatic recovery without explicitly enabling and disabling the port channel member ports at either end.

Table 14-1 compares On and Active modes.

Table 14-1 Channel Group Configuration Differences

On Mode	Active Mode
No protocol is exchanged.	A port channel protocol negotiation is performed with the peer ports.
Moves interfaces to the suspended state if its operational values are incompatible with the SAN port channel.	Moves interfaces to the isolated state if its operational values are incompatible with the SAN port channel.
When you add or modify a port channel member port configuration, you must explicitly disable (shut) and enable (no shut) the port channel member ports at either end.	When you add or modify a port channel interface, the SAN port channel automatically recovers.
Port initialization is not synchronized.	There is synchronized startup of all ports in a channel across peer switches.
All misconfigurations are not detected as no protocol is exchanged.	Consistently detect misconfigurations using a port channel protocol.
Transitions misconfigured ports to the suspended state. You must explicitly disable (shut) and enable (no shut) the member ports at either end.	Transitions misconfigured ports to the isolated state to correct the misconfiguration. Once you correct the misconfiguration, the protocol ensures automatic recovery.
This is the default mode.	You must explicitly configure this mode.

Send comments to nx5000-docfeedback@cisco.com

To configure active mode using Fabric Manager, perform this task:

-
- Step 1** Expand **ISLs**, and then choose **Port Channels** in the Physical Attributes pane.
You see the port channels configured in the Information pane.
- Step 2** Click the **Protocols** tab. From the Mode drop-down list, choose the appropriate mode for the Port Channel.
- Step 3** Click the **Apply Changes** icon to save any modifications.
-

About SAN Port Channel Deletion

When you delete the SAN port channel, the corresponding channel membership is also deleted. All interfaces in the deleted SAN port channel convert to individual physical links. After the SAN port channel is removed, regardless of the mode (active and on) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the [“Setting the Interface Administrative State”](#) section on page 10-9).

If you delete the SAN port channel for one port, then the individual ports within the deleted SAN port channel retain the compatibility parameter settings (speed, mode, port VSAN, allowed VSAN, and port security). You can explicitly change those settings as required.


- If you use the default On mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the Active mode, then the port channel ports automatically recover from the deletion.

Deleting SAN Port Channels

To delete a port channel using the Port Channel Wizard in Fabric Manager, perform this task:

-
- Step 1** Click the **Port Channel Wizard** icon in the toolbar (see [Figure 14-12](#)).

Figure 14-12 Port Channel Wizard Icon



You see the first Port Channel Wizard screen.

- Step 2** Select the existing port channel that you want to delete and click **Next**.
You see a list of the ISLs currently associated with this port channel.
- Step 3** Click **Next**.
You see an editable list of associated ISLs and available ISLs for this port channel.
- Step 4** Click each associated ISL and click the **left arrow** to remove all ISLs from the port channel.

Send comments to nx5000-docfeedback@cisco.com

- Step 5** Check the **Delete Port Channel If Empty** check box to delete this port channel.
- Step 6** Click **Finish** to save any modifications or click **Cancel** to discard any changes.
-

Interfaces in a SAN Port Channel

You can add or remove a physical Fibre Channel interface (or a range of interfaces) to an existing SAN port channel. The compatible parameters on the configuration are mapped to the SAN port channel. Adding an interface to a SAN port channel increases the channel size and bandwidth of the SAN port channel. Removing an interface from a SAN port channel decreases the channel size and bandwidth of the SAN port channel.



Note

Virtual Fibre Channel interfaces cannot be added to SAN port channels.

This section describes interface configuration for a SAN port channel and includes the following topics:

- [About Interface Addition to a SAN Port Channel, page 14-12](#)
- [Adding an Interface to a SAN Port Channel, page 14-13](#)
- [Forcing an Interface Addition, page 14-14](#)
- [About Interface Deletion from a SAN Port Channel, page 14-14](#)
- [Deleting an Interface from a SAN Port Channel, page 14-14](#)

About Interface Addition to a SAN Port Channel

You can add a physical interface (or a range of interfaces) to an existing SAN port channel. The compatible parameters on the configuration are mapped to the SAN port channel. Adding an interface to a SAN port channel increases the channel size and bandwidth of the SAN port channel.

After the members are added, regardless of the mode (Active and On) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down.

Compatibility Check

A compatibility check ensures that the same parameter settings are used in all physical ports in the channel. Otherwise, they cannot become part of a SAN port channel. The compatibility check is performed before a port is added to the SAN port channel.

The check ensures that the following parameters and settings match at both ends of a SAN port channel:

- Capability parameters (type of interface, Fibre Channel at both ends).
- Administrative compatibility parameters (speed, mode, port VSAN, allowed VSAN, and port security).
- Operational parameters (speed and remote switch's WWN).

Send comments to nx5000-docfeedback@cisco.com

A port addition procedure fails if the capability and administrative parameters in the remote switch are incompatible with the capability and administrative parameters in the local switch. If the compatibility check is successful, the interfaces are operational and the corresponding compatibility parameter settings apply to these interfaces.

Suspended and Isolated States

If the operational parameters are incompatible, the compatibility check fails and the interface is placed in a suspended or isolated state based on the configured mode:

- An interface enters the suspended state if the interface is configured in the On mode.
- An interface enters the isolated state if the interface is configured in the Active mode.

See the “Reason Codes” section on page 10-5.

Adding an Interface to a SAN Port Channel

To add an interface or range of interfaces to a SAN port channel using Fabric Manager, perform this task:

- Step 1** Expand **ISLs**, and then choose **Port Channels** in the Physical Attributes pane. You see the SAN port channels configured in the Information pane (see [Figure 14-13](#)).

Figure 14-13 Port Channels

Switch	Channel	Force	Members Admin	Members Oper	Last Status	Last FailureCause	Last Time
sw172-22-46-223	channel1	<input type="checkbox"/>	fcip4	fcip4	successful		2006/02/22-12:15:31
sw172-22-46-220	channel1	<input type="checkbox"/>	fcip5,fcip7,fcip8,fcip11	fcip5,fcip7,fcip8,fcip11	successful		2006/02/23-12:33:52
sw172-22-46-233	channel10	<input type="checkbox"/>	fcip5,fcip7,fcip8,fcip11	fcip5,fcip7,fcip8,fcip11	successful		2006/02/23-12:15:26
sw172-22-46-174	channel1	<input type="checkbox"/>	fcip5,fcip7,fcip8,fcip11	fcip5,fcip7,fcip8,fcip11	successful		2006/02/23-12:15:31
sw172-22-46-223	channel10	<input type="checkbox"/>	gggE2/1	gggE2/1	successful		2006/02/23-12:15:31
sw172-22-46-220	channel2	<input type="checkbox"/>	fcip6	fcip6	successful		2006/02/23-12:33:52
sw172-22-46-220	channel3	<input type="checkbox"/>	fcip4	fcip4	successful		2006/02/23-12:33:52
sw172-22-46-220	channel4	<input type="checkbox"/>	fcip4	fcip4	successful		2006/02/23-12:33:52
sw172-22-46-220	channel5	<input type="checkbox"/>	gggE9/5	gggE9/5	successful		2006/02/23-12:33:52
sw172-22-46-220	channel10	<input type="checkbox"/>	gggE9/5	gggE9/5	successful		2006/02/22-12:15:11

- Step 2** Click the **Channels** tab and find the switch and SAN port channel that you want to edit.
- Step 3** Set Members Admin to the interface or list of interfaces that you want to add to the SAN port channel.
- Step 4** Click the **Apply Changes** icon to save any modifications or click **Undo Changes** to discard any changes.

Send comments to nx5000-docfeedback@cisco.com

Forcing an Interface Addition

You can force the port configuration to be overwritten by the SAN port channel. In this case, the interface is added to a SAN port channel.

- If you use the default On mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the Active mode, then the port channel ports automatically recover from the addition.



Note

When SAN port channels are created from within an interface, the **force** option cannot be used.

After the members are forcefully added, regardless of the mode (Active and On) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the “[Setting the Interface Administrative State](#)” section on page 10-9).

To force the addition of a port to a SAN port channel using Fabric Manager, perform this task:

-
- Step 1** Expand **ISLs**, and then choose **Port Channels** in the Physical Attributes pane.
You see the port channels configured in the Information pane.
 - Step 2** Click the **Channels** tab and find the switch and SAN port channel that you want to edit.
 - Step 3** Set Members Admin to the interface or list of interfaces that you want to add to the SAN port channel.
 - Step 4** Check the **Force** check box to force this interface addition.
 - Step 5** Click the **Apply Changes** icon to save any modifications.
-

About Interface Deletion from a SAN Port Channel

When a physical interface is deleted from the SAN port channel, the channel membership is automatically updated. If the deleted interface is the last operational interface, then the port channel status is changed to a down state. Deleting an interface from a SAN port channel decreases the channel size and bandwidth of the SAN port channel.

- If you use the default On mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the Active mode, then the port channel ports automatically recover from the deletion.

After the members are deleted, regardless of the mode (Active and On) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down.

Deleting an Interface from a SAN Port Channel

To delete a physical interface (or a range of physical interfaces) from a SAN port channel using Fabric Manager, perform this task:

-
- Step 1** Expand **ISLs**, and then choose **Port Channels** in the Physical Attributes pane.
You see the SAN port channels configured in the Information pane.

Send comments to nx5000-docfeedback@cisco.com

- Step 2** Click the **Channels** tab and find the switch and SAN port channel that you want to edit.
 - Step 3** Remove the interface or list of interfaces that you want deleted in the Members the Admin column.
 - Step 4** Click the **Apply Changes** icon to save any modifications.
-

Port Channel Protocol

The switch software provides robust error detection and synchronization capabilities. You can manually configure channel groups, or they can be automatically created. In both cases, the channel groups have the same capability and configurational parameters. Any change in configuration applied to the associated port channel interface is propagated to all members of the channel group.

Cisco SAN switches support a protocol to exchange port channel configurations, which simplifies port channel management with incompatible ISLs. An additional autocreation mode enables ISLs with compatible parameters to automatically form channel groups without manual intervention.

The port channel protocol is enabled by default.

The port channel protocol expands the port channel functional model in Cisco SAN switches. It uses the exchange peer parameters (EPP) services to communicate across peer ports in an ISL. Each switch uses the information received from the peer ports along with its local configuration and operational values to decide if it should be part of a SAN port channel. The protocol ensures that a set of ports are eligible to be part of the same SAN port channel. They are only eligible to be part of the same port channel if all the ports have a compatible partner.

The port channel protocol uses two subprotocols:

- **Bringup protocol**—Automatically detects misconfigurations so you can correct them. This protocol synchronizes the SAN port channel at both ends so that all frames for a given flow (as identified by the source FC ID, destination FC ID and OX_ID) are carried over the same physical link in both directions. This helps make applications such as write acceleration work for SAN port channels over FCIP links.
- **Autocreation protocol**—Automatically aggregates compatible ports into a SAN port channel.

This section describes how to configure the port channel protocol and includes the following sections:

- [About Channel Group Creation, page 14-15](#)
- [Autocreation Guidelines, page 14-17](#)
- [Enabling and Configuring Autocreation, page 14-17](#)
- [About Manually Configured Channel Groups, page 14-18](#)
- [Converting to Manually Configured Channel Groups, page 14-18](#)

About Channel Group Creation


If channel group autocreation is enabled, ISLs can be configured automatically into channel groups without manual intervention. [Figure 14-14](#) shows an example of channel group autocreation.

The first ISL comes up as an individual link. In the example shown in [Figure 14-14](#), this is link A1-B1. When the next link comes up (A2-B2 in the example), the port channel protocol determines if this link is compatible with link A1-B1 and automatically creates channel groups 10 and 20 in the respective

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

switches. Link A3-B3 can join the channel groups (and the port channels) if the respective ports have compatible configurations. Link A4-B4 operates as an individual link, because it is not compatible with the existing member ports in the channel group.

Figure 14-14 Autocreating Channel Groups



The channel group numbers are assigned dynamically (when the channel group is formed).

The channel group number may change across reboots for the same set of port channels depending on the initialization order of the ports.

Table 14-2 identifies the differences between user-configured and auto-configured channel groups.

Table 14-2 Channel Group Configuration Differences

User-Configured Channel Group	Autocreated Channel Group
Manually configured by the user.	Created automatically when compatible links come up between two compatible switches, if channel group autocreation is enabled in all ports at both ends.
Member ports cannot participate in autocreation of channel groups. The autocreation feature cannot be configured.	None of these ports are members of a user-configured channel group.
You can form the SAN port channel with a subset of the ports in the channel group. Incompatible ports remain in a suspended or isolated state depending on the On or Active mode configuration.	All ports included in the channel group participate in the SAN port channel. No member port becomes isolated or suspended; instead, the member port is removed from the channel group when the link is found to be incompatible.
Any administrative configuration made to the SAN port channel is applied to all ports in the channel group, and you can save the configuration for the port channel interface.	Any administrative configuration made to the SAN port channel is applied to all ports in the channel group, but the configurations are saved for the member ports; no configuration is saved for the port channel interface. You can explicitly convert this channel group, if required.
You can remove any channel group and add members to a channel group.	You cannot remove a channel group. You cannot add members to the channel group or remove members. The channel group is removed when no member ports exist.

Send comments to nx5000-docfeedback@cisco.com

Autocreation Guidelines

When using the autocreation protocol, follow these guidelines:

- A port is not allowed to be configured as part of a SAN port channel when the autocreation feature is enabled. These two configurations are mutually exclusive.
- Autocreation must be enabled in both the local and peer ports to negotiate a SAN port channel.
- Aggregation occurs in one of two ways:
 - A port is aggregated into a compatible autocreated SAN port channel.
 - A port is aggregated with another compatible port to form a new SAN port channel.
- Newly created SAN port channels are allocated from the maximum possible port channel number in a decreasing order based on availability. If all port channel numbers are used up, aggregation is not allowed.
- You cannot change the membership or delete an autocreated SAN port channel.
- When you disable autocreation, all member ports are removed from the autocreated SAN port channel.
- Once the last member is removed from an autocreated SAN port channel, the channel is automatically deleted and the number is released for reuse.
- An autocreated SAN port channel is not persistent through a reboot. An autocreated SAN port channel can be manually configured to appear the same as a persistent SAN port channel. Once the SAN port channel is made persistent, the autocreation feature is disabled in all member ports.
- You can enable or disable the autocreation feature on a per-port basis or for all ports in the switch. When this configuration is enabled, the channel group mode is assumed to be active. The default for this task is disabled.
- If autocreation of channel groups is enabled for an interface, you must first disable autocreation before downgrading to earlier software versions or before configuring the interface in a manually configured channel group.



Tip

When enabling autocreation in any switch in the Cisco Nexus 5000 Series, we recommend that you retain at least one interconnected port between the switches without any autocreation configuration. If all ports between two switches are configured with the autocreation feature at the same time, a possible traffic disruption may occur between these two switches as ports are automatically disabled and reenabled when they are added to an autocreated SAN port channel.

Enabling and Configuring Autocreation

To configure port channel autocreation, check the **Dynamically form Port Channel Group from selected ISLs** option in the Port Channel Wizard. See the [“Configuring SAN Port Channels”](#) section on page 14-7.

Send comments to nx5000-docfeedback@cisco.com

About Manually Configured Channel Groups

A user-configured channel group cannot be converted to an autogenerated channel group. However, you can convert an autogenerated channel group to a manual channel group. This task is irreversible. The channel group number does not change, but the member ports operate according to the properties of the manually configured channel group, and channel group autocreation is implicitly disabled for all the member ports.



Tip

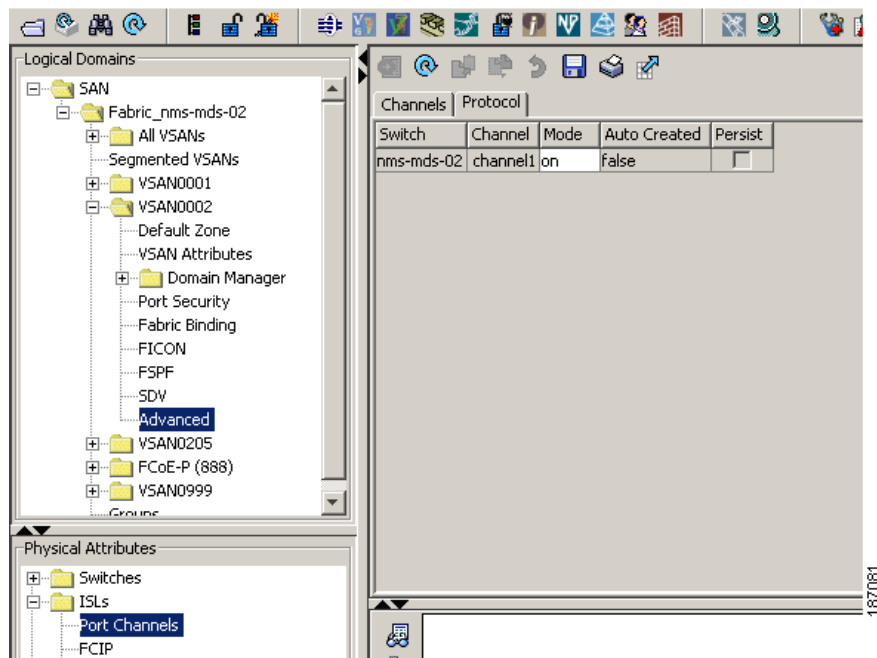
If you enable persistence, be sure to enable it at both ends of the SAN port channel.

Converting to Manually Configured Channel Groups

To convert an autogenerated channel group to a user-configured channel group using Fabric Manager, perform this task:

- Step 1** Expand **ISLs**, and then choose **Port Channels** in the Physical Attributes pane. Click the **Protocol** tab. You see the switch protocols as shown in [Figure 14-15](#).

Figure 14-15 Switch Protocols



- Step 2** Check the **Persist** check box for each channel that you want to convert to a manually configured channel group.
- Step 3** Click the **Apply Changes** icon to save any modifications.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Verifying SAN Port Channel Configuration

You can use the Information pane in Fabric Manager to verify your port channel configuration (see [Figure 14-16](#)).

Figure 14-16 Port Channel Summary in Fabric Manager

Switch	Channel	Force	Members Admin	Members Oper	Last Status	Last FailureCause	Last Time	CreationTime
sw172-22-46-220	channel1	<input type="checkbox"/>	Fcp0	Fcp0	successful		2007/04/17-21:05:20	2007/04/17-21:05:20
sw172-22-46-223	channel1	<input type="checkbox"/>	Fc1/1	Fc1/1	successful		2007/04/17-13:57:37	2007/04/17-13:57:37
sw172-22-46-221	channel1	<input checked="" type="checkbox"/>	Fc2/25	Fc2/25	successful		2007/04/17-14:50:03	2007/04/17-14:50:03
sw172-22-46-174	channel1	<input checked="" type="checkbox"/>	Fcp3	Fcp3	successful		2007/04/19-11:24:16	2007/04/19-11:24:15
sw172-22-46-220	channel2	<input type="checkbox"/>	Fc2/15	Fc2/15	successful		2007/04/17-21:05:20	2007/04/17-21:05:20
sw172-22-46-223	channel10	<input type="checkbox"/>	gg6/2/1	gg6/2/1	successful		2007/04/17-13:57:37	2007/04/17-13:57:37
sw172-22-46-174	channel2	<input checked="" type="checkbox"/>	Fcp4	Fcp4	successful		2007/04/19-11:24:16	2007/04/19-11:24:15
sw172-22-46-220	channel1	<input type="checkbox"/>	Fcp1	Fcp1	successful		2007/04/17-21:05:20	2007/04/17-21:05:20
sw172-22-46-220	channel1	<input type="checkbox"/>	Fc2/16	Fc2/16	successful		2007/04/17-21:05:20	2007/04/17-21:05:20
sw172-22-46-220	channel10	<input type="checkbox"/>	gg6/9/5	gg6/9/5	successful		2007/04/17-21:05:20	2007/04/17-21:05:20

Default Settings

[Table 14-3](#) lists the default settings for SAN port channels.

Table 14-3 Default SAN Port Channel Parameters

Parameters	Default
Port channels	FSPF is enabled by default.
Create port channel	Administratively up.
Default port channel mode	On.
Autocreation	Disabled.

Send comments to nx5000-docfeedback@cisco.com



Configuring and Managing VSANs

You can achieve higher security and greater stability in Fibre Channel fabrics by using virtual SANs (VSANs). VSANs provide isolation among devices that are physically connected to the same fabric. With VSANs you can create multiple logical SANs over a common physical infrastructure. Each VSAN can contain up to 239 switches and has an independent address space that allows identical Fibre Channel IDs (FC IDs) to be used simultaneously in different VSANs. This chapter includes the following sections:

- [Information About VSANs, page 15-1](#)
- [Configuring VSANs, page 15-5](#)
- [Default Settings, page 15-12](#)

Information About VSANs

A VSAN is a virtual storage area network (SAN). A SAN is a dedicated network that interconnects hosts and storage devices primarily to exchange SCSI traffic. In SANs you use the physical links to make these interconnections. A set of protocols run over the SAN to handle routing, naming, and zoning. You can design multiple SANs with different topologies.

This section describes VSANs and includes the following topics:

- [VSAN Topologies, page 15-1](#)
- [VSAN Advantages, page 15-4](#)
- [VSANs Versus Zones, page 15-4](#)

VSAN Topologies

With the introduction of VSANs, the network administrator can build a single topology containing switches, links, and one or more VSANs. Each VSAN in this topology has the same operation and property of a SAN. A VSAN has the following additional features:

- Multiple VSANs can share the same physical topology.
- The same Fibre Channel IDs (FC IDs) can be assigned to a host in another VSAN, which increases VSAN scalability.
- Every instance of a VSAN runs all required protocols such as FSPF, domain manager, and zoning.

Send comments to nx5000-docfeedback@cisco.com

- 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 15-1 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. No communication between VSANs is possible. Within each VSAN, all members can talk to one another.

Figure 15-1 Logical VSAN Segmentation





Figure 15-2 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 application servers or storage arrays can be connected to the switch using Fibre Channel or virtual Fibre Channel interfaces. A VSAN can include a mixture of Fibre Channel and virtual Fibre Channel interfaces.

Send comments to nx5000-docfeedback@cisco.com

Figure 15-2 Example of Two VSANs



The four switches in this network are interconnected by VSAN trunk links that carry both VSAN 2 and VSAN 7 traffic. You can configure a different inter-switch topology for each VSAN. In [Figure 15-2](#), the inter-switch topology is identical for VSAN 2 and VSAN 7.

Without VSANs, a network administrator would need separate switches and links for separate SANs. By enabling VSANs, the same switches and links may be shared by multiple VSANs. VSANs allow SANs to be built on port granularity instead of switch granularity. [Figure 15-2](#) illustrates that a VSAN is a group of hosts or storage devices that communicate with each other using a virtual topology defined on the physical SAN.

The criteria for creating such groups differ based on the VSAN topology:

- VSANs can separate traffic based on the following requirements:
 - Different customers in storage provider data centers
 - Production or test in an enterprise network
 - Low and high security requirements
 - Backup traffic on separate VSANs
 - Replicating data from user traffic
- VSANs can meet the needs of a particular department or application.

Send comments to nx5000-docfeedback@cisco.com

VSAN Advantages

VSANs offer the following advantages:

- Traffic isolation—Traffic is contained within VSAN boundaries and devices reside only in one VSAN ensuring absolute separation between user groups, if desired.
- Scalability—VSANs are overlaid on top of a single physical fabric. The ability to create several logical VSAN layers increases the scalability of the SAN.
- Per VSAN fabric services—Replication of fabric services on a per VSAN basis provides increased scalability and availability.
- Redundancy—Several VSANs created on the same physical SAN ensure redundancy. If one VSAN fails, redundant protection (to another VSAN in the same physical SAN) is configured using a backup path between the host and the device.
- Ease of configuration—Users can be added, moved, or changed between VSANs without changing the physical structure of a SAN. Moving a device from one VSAN to another only requires configuration at the port level, not at a physical level.

Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

VSANs Versus Zones

Zones are always contained within a VSAN. You can define multiple zones in a VSAN.

Because two VSANs are equivalent to two unconnected SANs, zone A on VSAN 1 is different and separate from zone A in VSAN 2. [Table 15-1](#) lists the differences between VSANs and zones.

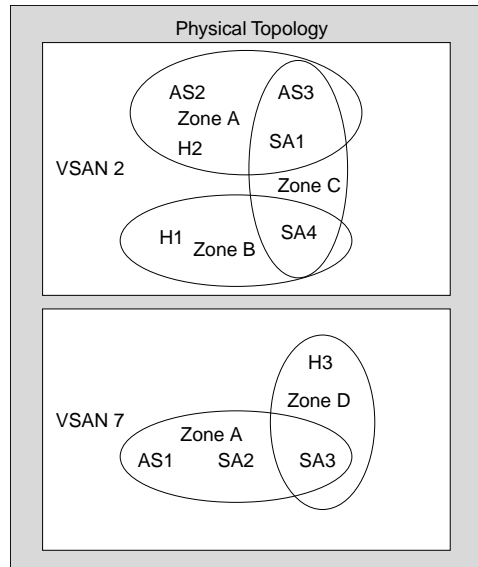
Table 15-1 VSAN and Zone Comparison

VSAN Characteristic	Zone Characteristic
VSANs equal SANs with routing, naming, and zoning protocols.	Routing, naming, and zoning protocols are not available on a per-zone basis.
VSANs limit unicast, multicast, and broadcast traffic.	Zones limit unicast traffic.
Membership is typically defined using the VSAN ID to F ports.	Membership is typically defined by the pWWN.
An HBA or a storage device can belong only to a single VSAN (the VSAN associated with the F 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 15-3](#) 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. Zone A defined in VSAN 2 is different and separate from zone A defined in VSAN 7.

Send comments to nx5000-docfeedback@cisco.com

Figure 15-3 VSANS with Zoning



Configuring VSANs

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.

This section describes how to create and configure VSANs and includes the following topics:

- [About VSAN Creation, page 15-6](#)

Send comments to nx5000-docfeedback@cisco.com

- [Creating VSANs Statically](#), page 15-6
- [About Port VSAN Membership](#), page 15-8
- [Assigning Static Port VSAN Membership](#), page 15-8
- [About the Default VSAN](#), page 15-8
- [About the Isolated VSAN](#), page 15-8
- [Displaying Isolated VSAN Membership](#), page 15-9
- [Operational State of a VSAN](#), page 15-9
- [About Static VSAN Deletion](#), page 15-9
- [Deleting Static VSANs](#), page 15-10
- [About Load Balancing](#), page 15-11
- [Configuring Load Balancing](#), page 15-11
- [About Interop Mode](#), page 15-12

About VSAN Creation

A VSAN is in the operational state if the VSAN is active and at least one port is up. This state indicates that traffic can pass through this VSAN. This state cannot be configured.


Creating VSANs Statically

You cannot configure any application-specific parameters for a VSAN before creating the VSAN.

To create and configure VSANs using Fabric Manager, perform this task:

-
- Step 1** Click the **Create VSAN** icon (see [Figure 15-4](#)).

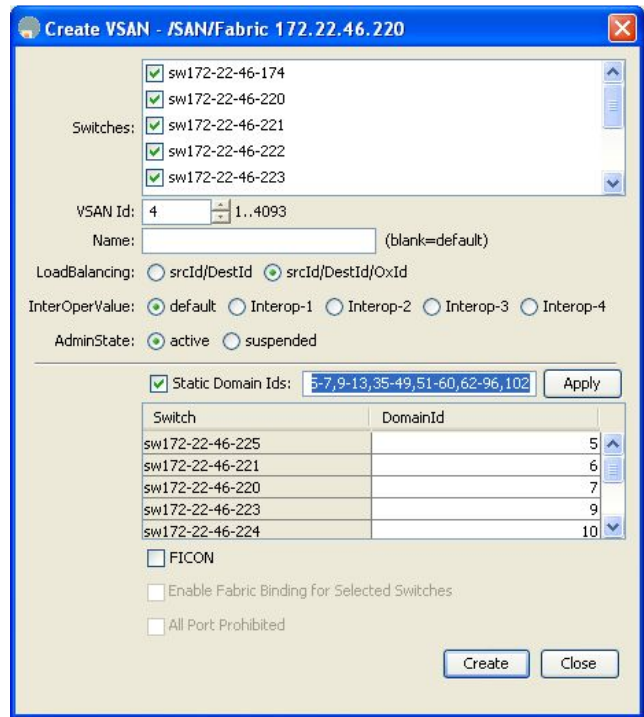
Figure 15-4 Create VSAN Icon



You see the Create VSAN dialog box as shown in [Figure 15-5](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 15-5 Create VSAN Dialog Box



Note If you check the Static Domain IDs check box, Fabric Manager creates the VSAN in suspended mode and then automatically activates the VSAN.

- Step 2** Check the switches that you want in this VSAN.
- Step 3** Fill in the VSAN Name and VSAN ID fields.
- Step 4** Set the LoadBalancing value and the InterOperValue.
- Step 5** Set the Admin State to **active** or **suspended**.
- Step 6** Check the **Static Domain Ids** check box to assign an unused static domain ID to the VSAN.
- Step 7** (Optional) Check the **Enable Fabric Binding for Selected Switches** options if you want this feature enabled.
See [Chapter 25, “Configuring Fabric Binding”](#) for details.
- Step 8** Complete the fields in this dialog box and click **Create** to add the VSAN or click **Close**.

Send comments to nx5000-docfeedback@cisco.com

About Port VSAN Membership

Port VSAN membership on the switch is assigned on a port-by-port basis. By default each port belongs to the default VSAN. You can assign VSAN membership to ports using one of two methods:

- Statically—Assigning VSANs to ports.
See the “[Assigning Static Port VSAN Membership](#)” section on page 15-8.
- Dynamically—Assigning VSANs based on the device WWN. This method is referred to as dynamic port VSAN membership (DPVM). Cisco Nexus 5000 Series switches do not support DPVM.

VSAN trunking ports have an associated list of VSANs that are part of an allowed list (see [Chapter 13, “Configuring VSAN Trunking”](#)).

Assigning Static Port VSAN Membership

To statically assign VSAN membership for an interface using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces** and then choose **FC Physical**.
You see the interface configuration in the Information pane.
- Step 2** Choose the **General** tab.
You see the Fibre Channel general physical information. Enter the new VSAN in the PortVSAN field.
- Step 3** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-

About the Default VSAN

The factory settings for switches in the Cisco Nexus 5000 Series have only the default VSAN 1 enabled. We recommend that you do not use VSAN 1 as your production environment VSAN. If no VSANs are configured, all devices in the fabric are considered part of the default VSAN. By default, all ports are assigned to the default VSAN.



Note VSAN 1 cannot be deleted, but it can be suspended.



Note Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

About the Isolated VSAN

VSAN 4094 is an isolated VSAN. When a VSAN is deleted, all nontrunking ports are transferred to the isolated VSAN to avoid an implicit transfer of ports to the default VSAN or to another configured VSAN. This action ensures that all ports in the deleted VSAN become isolated (disabled).

Send comments to nx5000-docfeedback@cisco.com



Note

When you configure a port in VSAN 4094 or move a port to VSAN 4094, that port is immediately isolated.



Caution

Do not use an isolated VSAN to configure ports.



Note

Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

Displaying Isolated VSAN Membership

To display interfaces that exist in the isolated VSAN using Fabric Manager, perform this task:

-
- Step 1** Expand **Fabricxx**, and then choose **All VSANs** in the Logical Domains pane.
You see the VSAN configuration in the Information pane.
- Step 2** Click the **Isolated Interfaces** tab.
You see the interfaces that are in the isolated VSAN.
-

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.

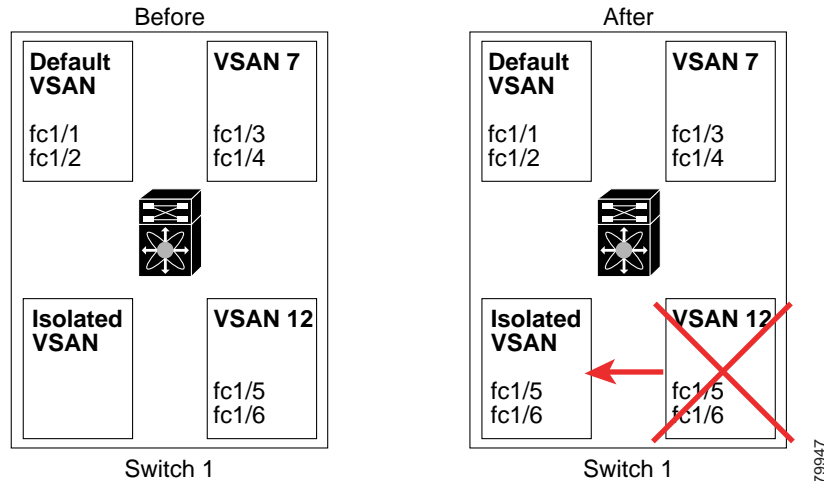
About Static VSAN Deletion

When an active VSAN is deleted, all of its attributes are removed from the running configuration. VSAN-related information is maintained by the system software as follows:

- VSAN attributes and port membership details are maintained by the VSAN manager. This feature is affected when you delete a VSAN from the configuration. When a VSAN is deleted, all the ports in that VSAN are made inactive and the ports are moved to the isolated VSAN. If the same VSAN is recreated, the ports do not automatically get assigned to that VSAN. You must explicitly reconfigure the port VSAN membership (see [Figure 15-6](#)).

Send comments to nx5000-docfeedback@cisco.com

Figure 15-6 VSAN Port Membership Details



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



Note

The allowed VSAN list is not affected when a VSAN is deleted (see [Chapter 13, “Configuring VSAN Trunking”](#)).

Any commands for a nonconfigured VSAN are rejected. For example, if VSAN 10 is not configured in the system, then a command request to move a port to VSAN 10 is rejected.

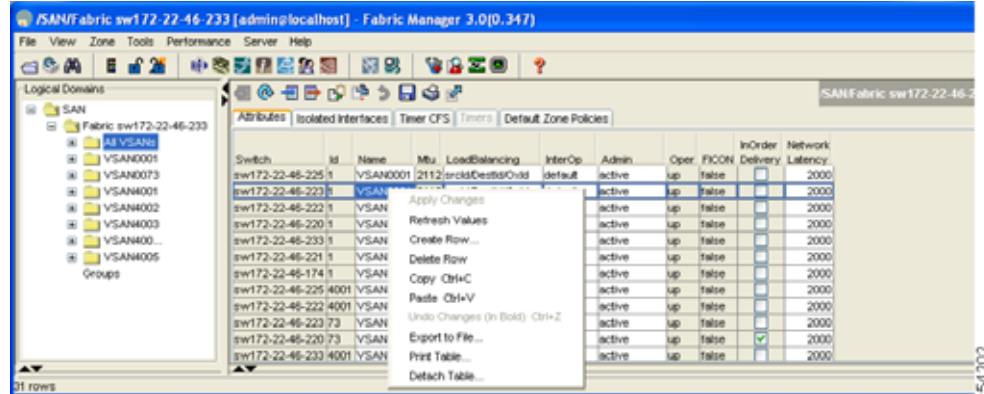
Deleting Static VSANs

To delete a VSAN and its attributes using Fabric Manager, perform this task:

- Step 1** Click **All VSANs** from the Logical Domains pane.
The VSANs in the fabric are listed in the Information pane.
- Step 2** Right-click the VSAN that you want to delete and click **Delete Row** from the drop-down menu (see [Figure 15-7](#)).

Send comments to nx5000-docfeedback@cisco.com

Figure 15-7 Deleting a VSAN



You see a confirmation dialog box.

- Step 3** Click **Yes** to confirm the deletion or **No** to close the dialog box without deleting the VSAN.

About Load Balancing

Load-balancing attributes indicate the use of the source-destination ID (src-dst-id) or the originator exchange OX ID (src-dst-ox-id, the default) for load-balancing path selection.

Configuring Load Balancing

To configure load balancing on an existing VSAN using Fabric Manager, perform this task:

- Step 1** Choose **Fabricxx > All VSANs** from the Logical Domains pane.

You see the VSAN configuration in the Information pane as shown in [Figure 15-8](#).

Figure 15-8 All VSAN Attributes



- Step 2** Click a VSAN and complete the LoadBalancing field.

Send comments to nx5000-docfeedback@cisco.com

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

About Interop Mode

Interoperability enables the products of multiple vendors to connect with each other. Fibre Channel standards guide vendors to create common external Fibre Channel interfaces. For additional information, see the “[Switch Interoperability](#)” section on page 22-7.

Default Settings

Table 15-2 lists the default settings for all configured VSANs.

Table 15-2 *Default VSAN Parameters*

Parameters	Default
Default VSAN	VSAN 1.
State	Active state.
Name	Concatenation of VSAN and a four-digit string representing the VSAN ID. For example, VSAN 3 is VSAN0003.
Load-balancing attribute	OX ID (src-dst-ox-id).



Configuring and Managing Zones

Zoning enables you to set up access control between storage devices or user groups. If you have administrator privileges in your fabric, you can create zones to increase network security and to prevent data loss or corruption. Zoning is enforced by examining the source-destination ID field.

Advanced zoning capabilities specified in the FC-GS-4 and FC-SW-3 standards are supported. You can use either the existing basic zoning capabilities or the advanced, standards-compliant zoning capabilities.

This chapter includes the following sections:

- [Information About Zoning, page 16-1](#)
- [Configuring Zones, page 16-7](#)
- [Zone Sets, page 16-12](#)
- [Zone Set Distribution, page 16-21](#)
- [Zone Set Duplication, page 16-24](#)
- [Verifying Zone Information, page 16-28](#)
- [Enhanced Zoning, page 16-29](#)
- [Compacting the Zone Database, page 16-33](#)
- [Default Settings, page 16-33](#)



Note

[Table 15-1 on page 15-4](#) lists the differences between zones and VSANs.

Information About Zoning

Zoning is described in the following topics:

- [Zoning Features, page 16-2](#)
- [Zoning Example, page 16-3](#)
- [Zone Implementation, page 16-4](#)
- [Active and Full Zone Set Configuration Guidelines, page 16-5](#)

Send comments to nx5000-docfeedback@cisco.com

Zoning Features

Zoning includes 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 physical fabric can have a maximum of 16,000 members. This includes all VSANs in the fabric.
- 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.
 - A zone switch can have a maximum of 500 zone sets.
- Zoning can be administered from any switch in the fabric.
 - When you activate a zone (from any switch), all switches in the fabric receive the active zone set. Additionally, full zone sets are distributed to all switches in the fabric, if this feature is enabled in the source switch.
 - If a new switch is added to an existing fabric, zone sets are acquired by the new switch.
- Zone changes can be configured nondisruptively.
 - New zones and zone sets can be activated without interrupting traffic on unaffected ports or devices.
- Zone membership can be specified using the following identifiers:
 - 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 a Cisco switch domain and additionally specifies a port belonging to a non-Cisco switch.



Note

For N ports attached to the switch over a virtual Fibre Channel interface, you can specify zone membership using the pWWN of the N port, the FC ID of the N port, or the fabric pWWN of the virtual Fibre Channel interface.

Send comments to nx5000-docfeedback@cisco.com

- Default zone membership includes all ports or WWNs that do not have a specific membership association. Access between default zone members is controlled by the default zone policy.
- You can configure up to 8000 zones per VSAN and a maximum of 8000 zones for all VSANs on the switch.




Note

Interface-based zoning only works with Cisco SAN switches. Interface-based zoning does not work for VSANs configured in interop mode.

Zoning Example

Figure 16-1 shows 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. H3 resides in both zones.


Figure 16-1 Fabric with Two Zones



You can use other ways to partition this fabric into zones. Figure 16-2 shows 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 only H2 and S2 in zone 3, and to H1 and S1 in zone 1.

Send comments to nx5000-docfeedback@cisco.com

Figure 16-2 Fabric with Three Zones



Zone Implementation

Cisco Nexus 5000 Series switches automatically support the following basic zone features (no additional configuration is required):

- Zones are contained in a VSAN.
- Hard zoning cannot be disabled.
- Name server queries are soft-zoned.
- Only active zone sets are distributed.
- Unzoned devices cannot access each other.
- A zone or zone set with the same name can exist in each VSAN.
- Each VSAN has a full database and an active database.
- Active zone sets cannot be changed, without activating a full zone database.
- Active zone sets are preserved across switch reboots.
- Changes to the full database must be explicitly saved.
- Zone reactivation (a zone set is active and you activate another zone set) does not disrupt existing traffic.

If required, you can additionally configure the following zone features:

- Propagate full zone sets to all switches on a per VSAN basis.
- Change the default policy for unzoned members.
- Interoperate with other vendors by configuring a VSAN in the interop mode. You can also configure one VSAN in the interop mode and another VSAN in the basic mode in the same switch without disrupting each other.
- Bring E ports out of isolation.

Send comments to nx5000-docfeedback@cisco.com

Active and Full Zone Set Configuration Guidelines

Before configuring a zone set, consider the following guidelines:

- Each VSAN can have multiple zone sets but only one zone set can be active at any given time.
- When you create a zone set, that zone set becomes a part of the full zone set.
- When you activate a zone set, a copy of the zone set from the full zone set is used to enforce zoning, and is called the active zone set. An active zone set cannot be modified. A zone that is part of an active zone set is called an active zone.
- The administrator can modify the full zone set even if a zone set with the same name is active. However, the modification will be enforced only upon reactivation.
- When the activation is done, the active zone set is automatically stored in persistent configuration. This enables the switch to preserve the active zone set information across switch resets.
- All other switches in the fabric receive the active zone set so they can enforce zoning in their respective switches.
- Hard and soft zoning are implemented using the active zone set. Modifications take effect during zone set activation.
- An FC ID or Nx port that is not part of the active zone set belongs to the default zone and the default zone information is not distributed to other switches.



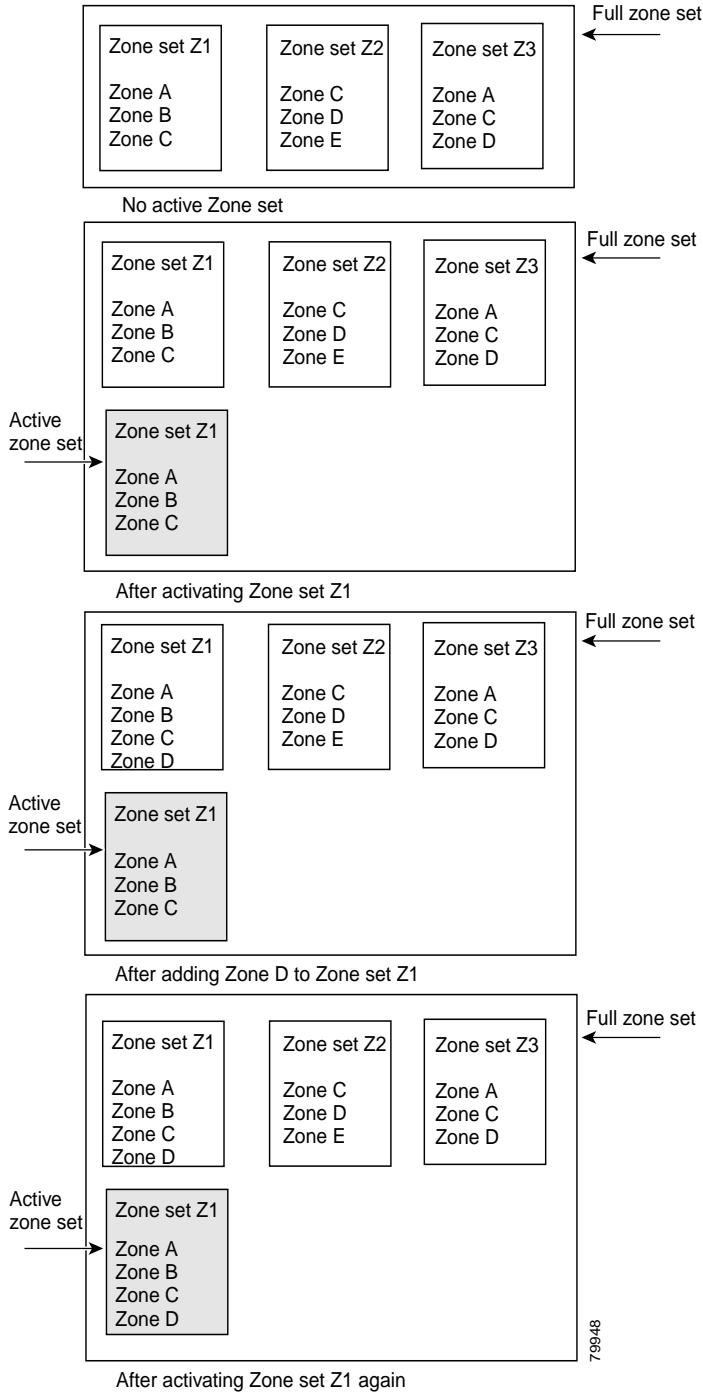
Note

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

[Figure 16-3](#) shows a zone being added to an activated zone set.

Send comments to nx5000-docfeedback@cisco.com

Figure 16-3 Active and Full Zone Sets



Send comments to nx5000-docfeedback@cisco.com

Configuring Zones

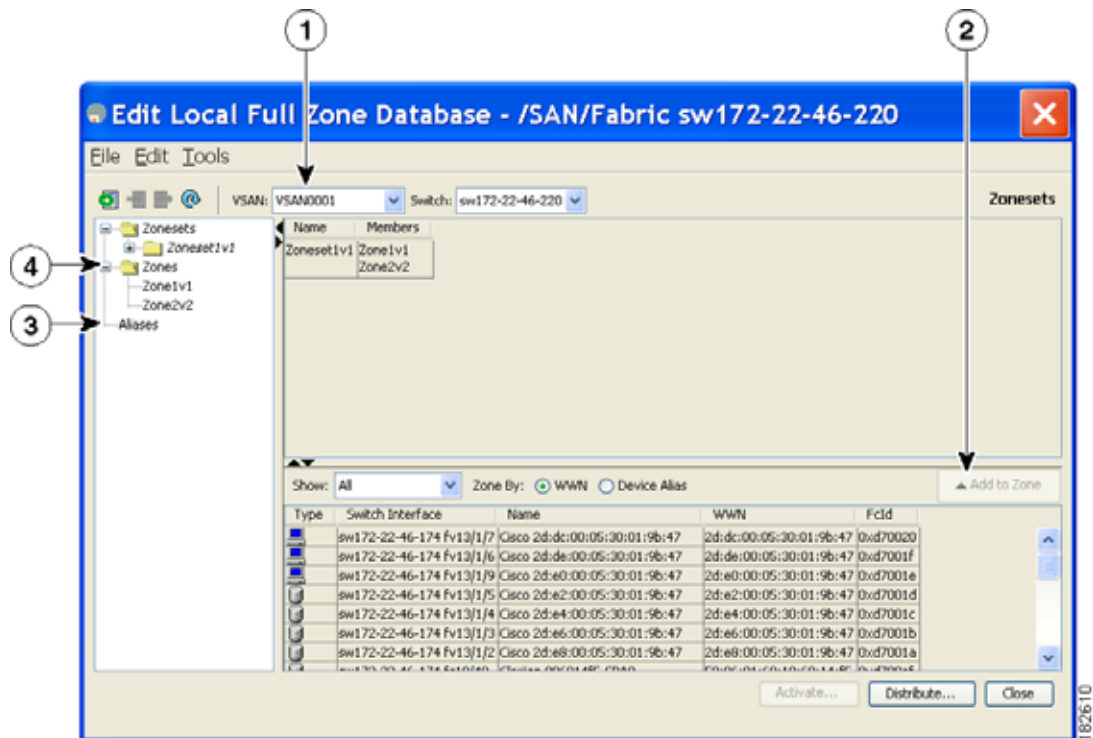
This section describes how to configure zones and includes the following topics:

- [About the Zone Configuration Tool, page 16-7](#)
- [Configuring Zones Using the Zone Configuration Tool, page 16-8](#)
- [Adding Zone Members, page 16-10](#)
- [Configuring the Default Zone Policy, page 16-12](#)

About the Zone Configuration Tool

The Zone Configuration tool allows you to zone across multiple switches and all zoning features are available through the Edit Local Full Zone Database dialog box (see [Figure 16-4](#)).

Figure 16-4 Edit Local Full Zone Database Dialog Box



1	You can display information for a specific VSAN by selecting the VSAN in the drop-down menu, and then pressing the Enter key.	3	You can add zoning characteristics based on the aliases in different folders.
2	You can use the Add to zone button to move devices up or down by alias or by zone.	4	You can triple-click to rename zone sets, zones, or aliases in the tree.

Send comments to nx5000-docfeedback@cisco.com

**Note**

The Device Alias radio button is visible only if device alias is in enhanced mode. For more information, see the “[Configuring Device Alias Modes](#)” section on page 17-4.

**Tip**


Expand **Switches** from the Physical Attributes pane to retrieve the switch world-wide name (sWWN). If you do not provide a sWWN, the software automatically uses the local sWWN.

Configuring Zones Using the Zone Configuration Tool

To create a zone and move it into a zone set using Fabric Manager, perform this task:

- Step 1** Click the **Zone** icon in the toolbar (See [Figure 16-5](#)).

Figure 16-5 Zone Icon



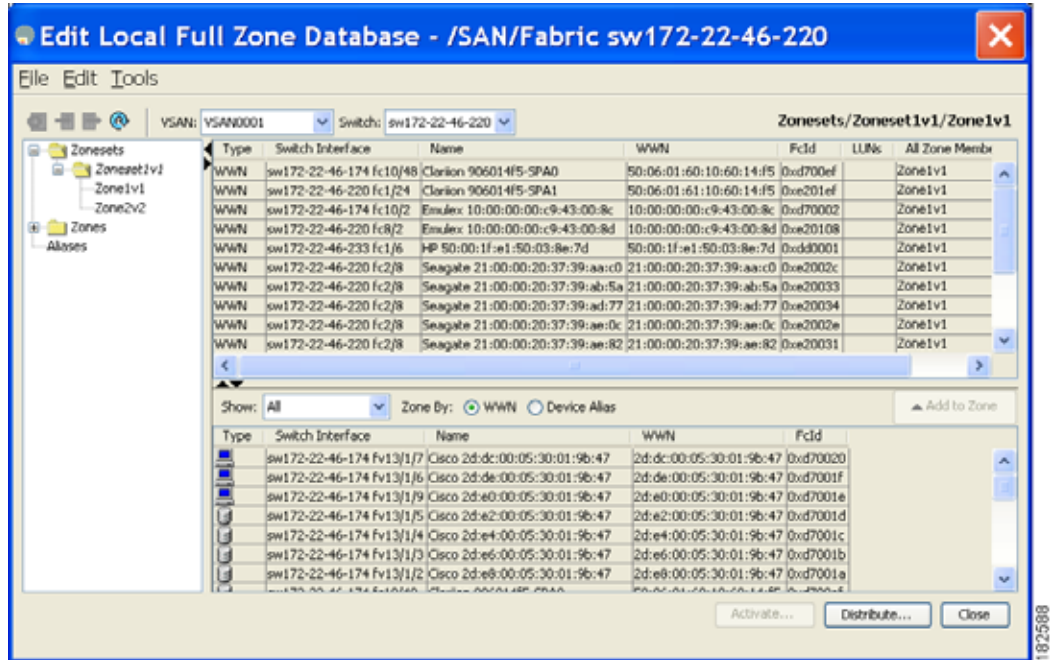
You see the Select VSAN dialog box.

- Step 2** Choose the VSAN where you want to create a zone and click **OK**.

You see the Edit Local Full Zone Database dialog box as shown in [Figure 16-6](#).

Send comments to nx5000-docfeedback@cisco.com


Figure 16-6 Edit Local Full Zone Database Dialog Box



If you want to view zone membership information, right-click in the **All Zone Membership(s)** column, and then choose **Show Details** for the current row or all rows from the pop-up menu.

- Step 3** Click **Zones** in the left pane and click the **Insert** icon to create a zone.
You see the Create Zone dialog box as shown in [Figure 16-7](#).

Figure 16-7 Create Zone Dialog Box



- Step 4** Enter a zone name.
- Step 5** Check one of the following check boxes:
- Read Only**—The zone permits read and denies write.
 - Permit QoS traffic with Priority**—You set the priority from the drop-down list.
 - Restrict Broadcast frames to Zone Members**
- Step 6** Click **OK** to create the zone.
If you want to move this zone into an existing zone set, skip to [Step 8](#).
- Step 7** Click **Zoneset** in the left pane and click the **Insert** icon to create a zone set.
You see the Zoneset Name dialog box as shown in [Figure 16-8](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 16-8 Zoneset Name Dialog Box



Step 8 Enter a zone set name and click **OK**.

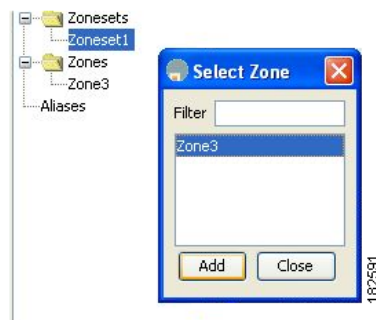


Note One of these symbols (\$, -, ^, _) or all alphanumeric characters are supported. In interop mode 2 and 3, this symbol (_) or all alphanumeric characters are supported.

Step 9 Choose the zone set where you want to add a zone and click the **Insert** icon, or you can drag and drop Zone3 over Zoneset1.

You see the Select Zone dialog box as shown in [Figure 16-9](#).

Figure 16-9 Select Zone Dialog Box



Step 10 Click **Add** to add the zone.

Adding Zone Members

After you create a zone, you can add members to the zone. You can add members using multiple port identification types.

To add a member to a zone using Fabric Manager, perform this task:

Step 1 Choose **Zone > Edit Local Full Zone Database**.

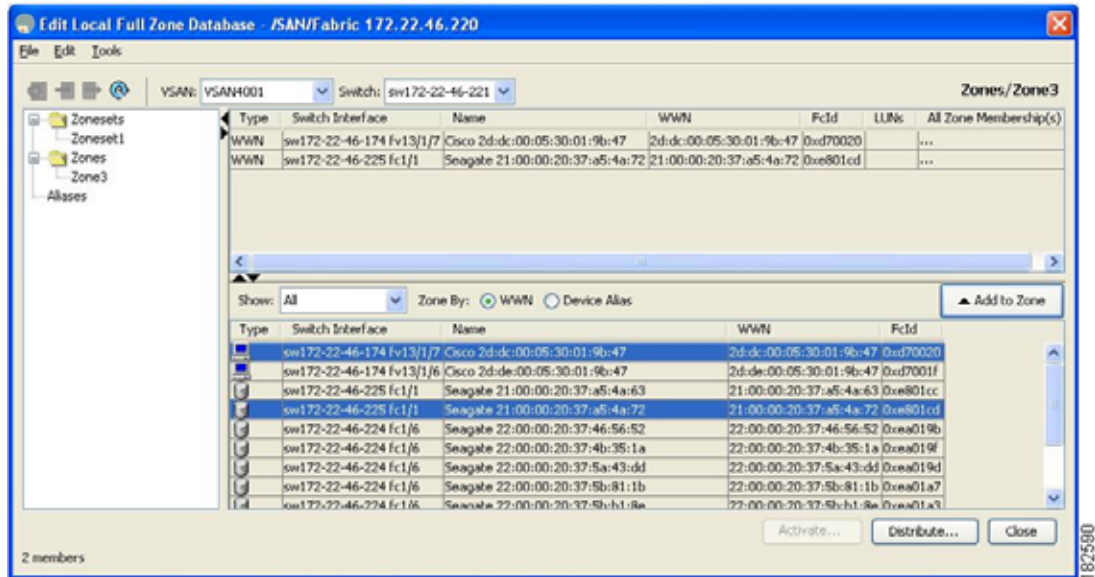
You see the Select VSAN dialog box.

Step 2 Choose a VSAN and click **OK**.

You see the Edit Local Full Zone Database dialog box (see [Figure 16-10](#)) for the selected VSAN.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

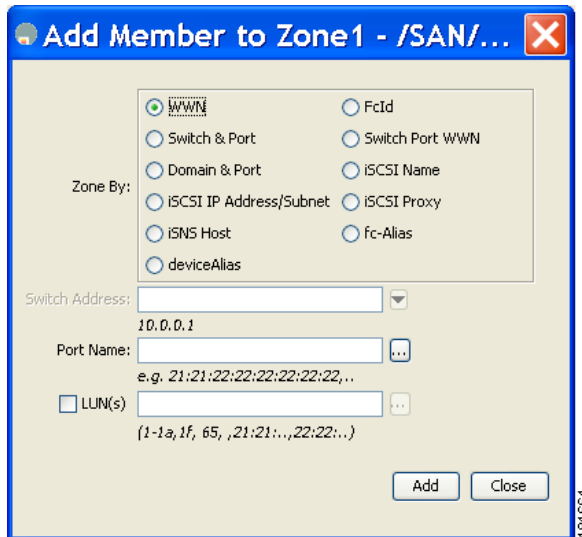
Figure 16-10 Edit Local Full Zone Database Dialog Box



Step 3 Select the members you want to add from the Fabric pane and click **Add to Zone** or click the zone where you want to add members and click the **Insert** icon.

You see the Add Member to Zone dialog box as shown in Figure 16-11.

Figure 16-11 Add Member to Zone Dialog Box



Note The Device Alias radio button is visible only if device alias is in enhanced mode. For more information, see the “[Configuring Device Alias Modes](#)” section on page 17-4.

Step 4 Click the browse button and choose a port name or check the **LUN** check box and click the browse button to configure LUNs.

Send comments to nx5000-docfeedback@cisco.com

Step 5 Click **Add** to add the member to the zone.



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

Configuring the Default Zone Policy

To permit or deny traffic in the default zone using Fabric Manager, perform this task:

Step 1 Choose **Zone > Edit Local Full Zone Database**.

You see the Select VSAN dialog box.

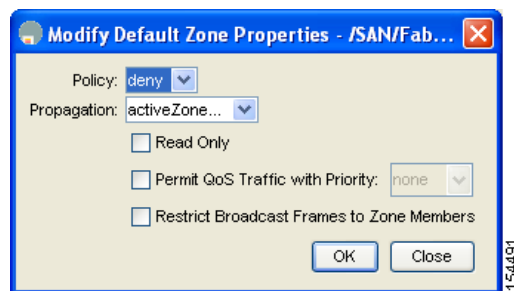
Step 2 Choose a VSAN and click **OK**.

You see the Edit Local Full Zone Database dialog box for the selected VSAN.

Step 3 Choose **Edit > Edit Default Zone Attributes** to configure the default zone QoS priority attributes.

You see the Modify Default Zone Properties dialog box as shown in [Figure 16-12](#).

Figure 16-12 Modify Default Zone Properties Dialog Box



Step 4 Set the Policy drop-down list to **permit** to permit traffic in the default zone, or set it to **deny** to block traffic in the default zone.

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

Zone Sets

This section describes zone sets and includes the following topics:

- [About Zone Set Creation, page 16-13](#)
- [Activating a Zone Set, page 16-13](#)
- [Displaying Zone Membership Information, page 16-15](#)
- [About the Default Zone, page 16-16](#)
- [Configuring the Default Zone, page 16-16](#)


Send comments to nx5000-docfeedback@cisco.com

- [About FC Alias Creation, page 16-17](#)
- [Creating FC Aliases, page 16-17](#)
- [Adding Members to Aliases, page 16-18](#)
- [Converting Zone Members to pWWN-Based Members, page 16-20](#)
- [Zone Enforcement, page 16-21](#)

About Zone Set Creation

In [Figure 16-13](#), two separate sets are created, each with its own membership hierarchy and zone members.

Figure 16-13 Hierarchy of Zone Sets, Zones, and Zone Members



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



Tip

Zone sets are configured with the names of the member zones and the VSAN (if the zone set is in a configured VSAN).

Activating a Zone Set

Changes to a zone set do not take effect in a full zone set until you activate it.

To activate an existing zone set using Fabric Manager, perform this task:

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

Send comments to nx5000-docfeedback@cisco.com


Step 2 Choose a VSAN and click **OK**.

You see the Edit Local Full Zone Database dialog box for the selected VSAN.

Step 3 Click **Activate** to activate the zone set.

You see the preactivation check dialog box as shown in [Figure 16-14](#).

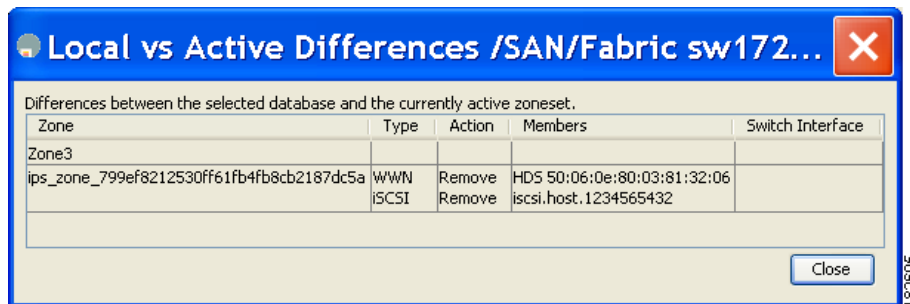
Figure 16-14 Pre-Activation Check Dialog Box



Step 4 Click **Yes** to review the differences.

You see the Local vs. Active Differences dialog box as shown in [Figure 16-15](#).


Figure 16-15 Local vs. Active Differences Dialog Box



Step 5 Click **Close** to close the dialog box.

You see the Save Configuration dialog box as shown in [Figure 16-16](#).

Figure 16-16 Save Configuration Dialog Box



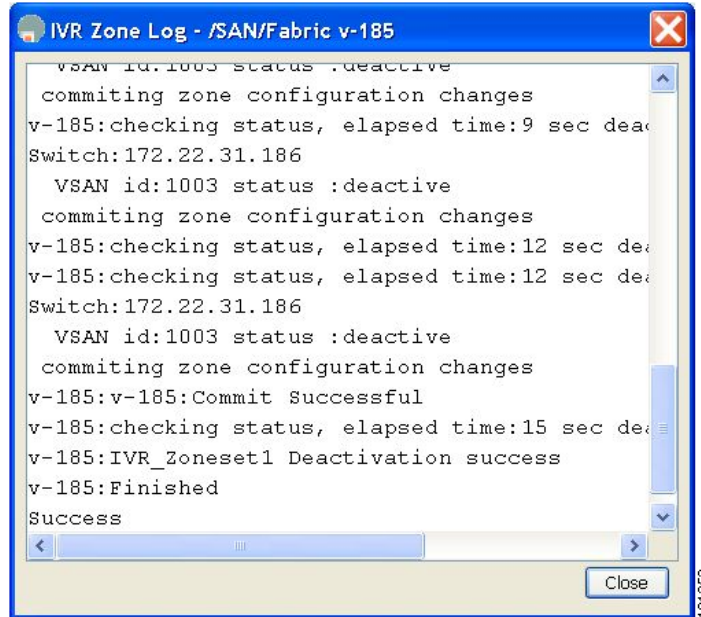
Step 6 Check the **Save Running to Startup Configuration** check box to save all changes to the startup configuration.

Step 7 Click **Continue Activation** to activate the zone set, or click **Cancel** to close the dialog box and discard any unsaved changes.

Send comments to nx5000-docfeedback@cisco.com

You see the Zone Log dialog box, which shows if the zone set activation was successful (see Figure 16-17).

Figure 16-17 Zone Log Dialog Box



To deactivate an existing zone set, perform this task:

- Step 1 Right-click the zone set you want to deactivate, and then choose **Deactivate** from the pop-up menu.
- Step 2 Click **OK** in the confirmation dialog box to deactivate the zone set.

Displaying Zone Membership Information

To display zone membership information for members assigned to zones in Fabric Manager, perform this task:

- Step 1 Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2 Choose a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
- Step 3 Click **Zones** in the left pane. The right pane lists the members for each zone.

Send comments to nx5000-docfeedback@cisco.com



Note The default zone members are explicitly listed only when the default zone policy is configured as **permit**. When the default zone policy is configured as **deny**, the members of this zone are not shown. See the “[Verifying Zone Information](#)” section on page 16-28.

About the Default Zone

Each member of a fabric (in effect a device attached to an Nx port) can belong to any zone. If a member is not part of any active zone, it is considered to be part of the default zone. Therefore, if no zone set is active in the fabric, all devices are considered to be in the default zone. Even though a member can belong to multiple zones, a member that is part of the default zone cannot be part of any other zone. The switch determines whether a port is a member of the default zone when the attached port comes up.



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

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



Note When the switch is initialized for the first time, no zones are configured and all members are considered to be part of the default zone. Members are not permitted to communicate with each other.

Configure the default zone policy on each switch in the fabric. If you change the default zone policy on one switch in a fabric, be sure to change it on all the other switches in the fabric.



Note The default settings for default zone configurations can be changed.

The default zone members are explicitly listed when the default policy is configured as permit or when a zone set is active. When the default policy is configured as deny, the members of this zone are not explicitly enumerated when you view the active zone set.

You can change the default zone policy for any VSAN by choosing **VSANxx > Default Zone** from the Logical Domains pane and clicking the **Policies** tab. It is recommended that you establish connectivity among devices by assigning them to a nondefault zone.

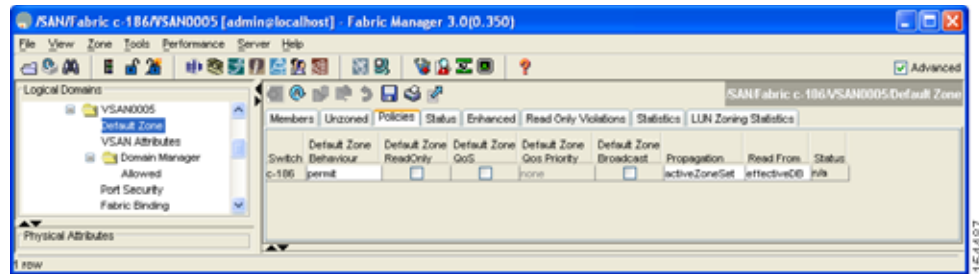
Configuring the Default Zone

To permit or deny traffic to members in the default zone using Fabric Manager, perform this task:

-
- Step 1** Expand a **VSAN**, and then choose **Default Zone** in the Fabric Manager Logical Domains pane.
 - Step 2** Click the **Policies** tab in the Information pane.
You see the zone policies information in the Information pane (see [Figure 16-18](#)).

Send comments to nx5000-docfeedback@cisco.com

Figure 16-18 Default Zone Policies



The active zone set is shown in italic type. After you make changes to the active zone set and before you activate the changes, the zone set is shown in boldface italic type.

- Step 3** In the Default Zone Behavior field, choose either **permit** or **deny** from the drop-down list.

About FC Alias Creation

You can assign an alias name and configure an alias member using the following values:

- pWWN—The WWN of the N port is in hex format (for example, 10:00:00:23:45:67:89:ab).
- fWWN—The WWN of the fabric port name is in hex format (for example, 10:00:00:23:45:67:89:ab).
- FC ID—The N port ID is in 0xhhhhhh format (for example, 0xce00d1).
- 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.
- Interface—Interface-based zoning is similar to port-based zoning because the switch interface is used to configure the zone. You can specify a switch interface as a zone member for both local and remote switches. To specify a remote switch, enter the remote switch WWN (sWWN) or the domain ID in the particular VSAN.



Tip

The switch supports a maximum of 2048 aliases per VSAN.


Creating FC Aliases

To create an FC alias using Fabric Manager, perform this task:

- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Choose a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
- Step 3** Click **FC-Aliases** in the lower left pane (see [Figure 16-19](#)). The right pane lists the existing aliases.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)


Figure 16-19 Creating an FC Alias



Step 4 Click the **Insert** icon to create an alias.

You see the Create Alias dialog box as shown in [Figure 16-20](#).

Figure 16-20 Create Alias Dialog Box



Step 5 Set the Alias Name and the pWWN.

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

Adding Members to Aliases

To add a member to an alias using Fabric Manager, perform this task:

Step 1 Choose **Zone > Edit Local Full Zone Database**.

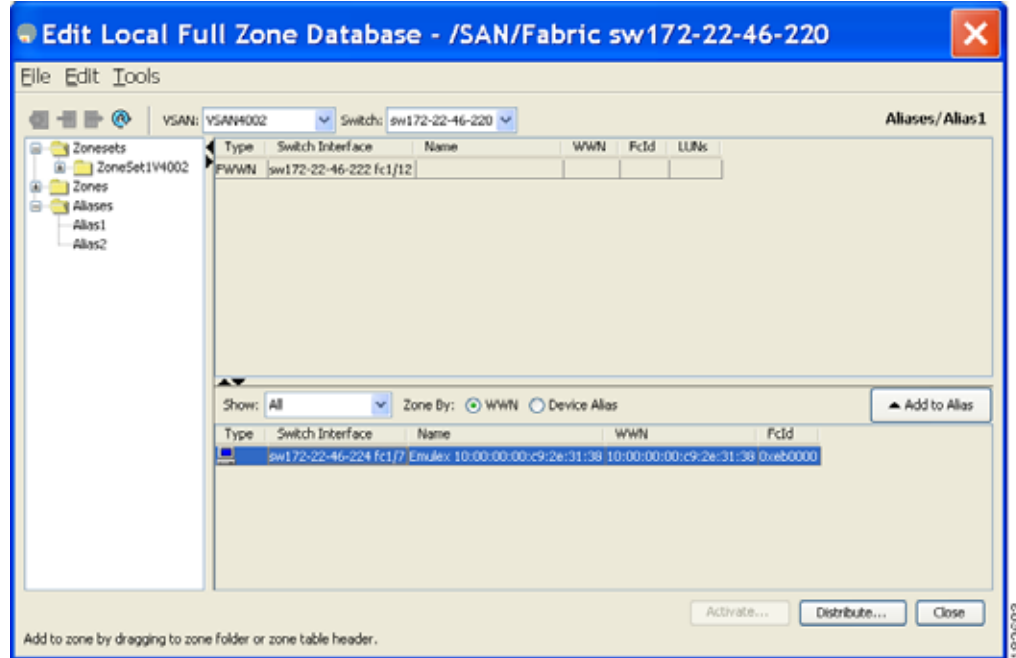
You see the Select VSAN dialog box.

Step 2 Choose a VSAN and click **OK**.

You see the Edit Local Full Zone Database dialog box for the selected VSAN as shown in [Figure 16-21](#).

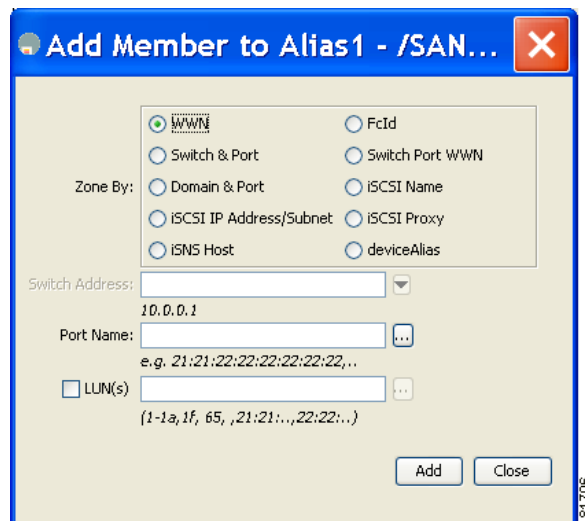
Send comments to nx5000-docfeedback@cisco.com

Figure 16-21 Edit Local Full Zone Database Dialog Box



- Step 3** Click the alias (in the FC-Aliases folder) where you want to add members.
- Step 4** Select the member(s) you want to add from the Fabric pane (see [Figure 16-21](#)) and click **Add to Alias**.
Fabric Manager provides an alternative method for adding members to the alias. Instead of step 4, perform steps 5 through 7.
- Step 5** Click the alias where you want to add members and click the **Insert** icon.
You see the Add Member to Alias dialog box as shown in [Figure 16-22](#).

Figure 16-22 Add Member to Alias Dialog Box



Send comments to nx5000-docfeedback@cisco.com



Note The Device Alias radio button is visible only if device alias is in enhanced mode. For more information, see [Configuring Device Alias Modes, page 17-4](#).

- Step 6** Click the browse button and choose a port name or check the **LUN** check box and click the browse button to configure LUNS.
- Step 7** Click **Add** to add the member to the alias.

Converting Zone Members to pWWN-Based Members

You can convert zone and alias members from switch port or FC ID based membership to pWWN-based membership. You can use this feature to convert to pWWN so that your zone configuration does not change if a switch is changed in your fabric.

To convert switch port and FC ID members to pWWN members using Fabric Manager, perform this task:

- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Choose a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
- Step 3** Click the zone you want to convert.
- Step 4** Choose **Tools > Convert Switch Port/FCID members to By pWWN**.
You see the conversion dialog box, listing all members that will be converted.
- Step 5** Verify the changes and click **Continue Conversion**.
- Step 6** Click **Yes** in the confirmation dialog box to convert that member to pWWN-based membership.



Tip You do not have to copy the running configuration to the startup configuration to store the active zone set. However, you need to copy the running configuration to the startup configuration to explicitly store full zone sets.



Note The pWWN of the virtual target does not appear in the zoning end devices database in Fabric Manager. If you want to zone the virtual device with a pWWN, you must enter it in the Add Member to Zone dialog box when creating a zone. However, if the device alias is in enhanced mode, the virtual device names appear in the device alias database in the Fabric Manager zoning window. In this case, users can choose to select either the device alias name or enter the pWWN in the Add Member to Zone dialog box. For more information, see the [“Adding Zone Members” section on page 16-10](#).



Note Be sure you understand how device alias modes work before enabling them. See [Chapter 17, “Distributing Device Alias Services”](#) for details and requirements about device alias modes.

Send comments to nx5000-docfeedback@cisco.com

Zone Enforcement

Zoning can be enforced in two ways: soft and hard. Each end device (N 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 N 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 N port. As frames enter the switch, source-destination IDs are compared with permitted combinations to allow the frame at wire speed. Hard zoning is applied to all forms of zoning.



Note

Hard zoning enforces zoning restrictions on every frame, and prevents unauthorized access.

Cisco Nexus 5000 Series switches support both hard and soft zoning.

Zone Set Distribution

You can distribute full zone sets using one of two methods: one-time distribution or full zone set distribution. [Table 16-1](#) lists the differences between the methods.

Table 16-1 Zone Set Distribution Differences

One-Time Distribution	Full Zone Set Distribution
Distributes the full zone set immediately.	Does not distribute the full zone set immediately.
Does not distribute the full zone set information along with the active zone set during activation, deactivation, or merge process.	Remembers to distribute the full zone set information along with the active zone set during activation, deactivation, and merge processes.

This section describes zone set distribution and includes the following topics:

- [Enabling Full Zone Set Distribution, page 16-21](#)
- [Enabling a One-Time Distribution, page 16-22](#)
- [About Recovering from Link Isolation, page 16-23](#)
- [Importing and Exporting Zone Sets, page 16-23](#)

Enabling Full Zone Set Distribution

All switches in the Cisco Nexus 5000 Series distribute active zone sets when new E port links come up or when a new zone set is activated in a VSAN. The zone set distribution takes effect while sending merge requests to the adjacent switch or while activating a zone set.

Send comments to nx5000-docfeedback@cisco.com

To enable full zone set and active zone set distribution to all switches on a per VSAN basis using Fabric Manager, perform this task:

-
- Step 1** Expand a **VSAN** and choose a zone set in the Logical Domains pane.
You see the zone set configuration in the Information pane. The Active Zones tab is the default.
- Step 2** Click the **Policies** tab.
You see the configured policies for the zone as shown in [Figure 16-23](#).

Figure 16-23 Configured Policies for the Zone

Switch	Default Zone Behaviour	Default Zone ReadOnly	Default Zone GoS	Default Zone QoS Priority	Default Zone Broadcast	Propagation	Read From	Status
sw172-22-46-182	deny	<input type="checkbox"/>	<input type="checkbox"/>	none	<input type="checkbox"/>	activeZoneSet	effectiveDB	file
sw172-22-46-224	deny	<input type="checkbox"/>	<input type="checkbox"/>	none	<input type="checkbox"/>	activeZoneSet	effectiveDB	file
sw172-22-46-221	deny	<input type="checkbox"/>	<input type="checkbox"/>	none	<input type="checkbox"/>	activeZoneSet	effectiveDB	file
sw172-22-46-223	deny	<input type="checkbox"/>	<input type="checkbox"/>	none	<input type="checkbox"/>	activeZoneSet	effectiveDB	file
sw172-22-46-220	deny	<input type="checkbox"/>	<input type="checkbox"/>	none	<input type="checkbox"/>	activeZoneSet	effectiveDB	file
sw172-22-46-233	deny	<input type="checkbox"/>	<input type="checkbox"/>	none	<input type="checkbox"/>	activeZoneSet	effectiveDB	file
sw172-22-46-225	deny	<input type="checkbox"/>	<input type="checkbox"/>	none	<input type="checkbox"/>	activeZoneSet	effectiveDB	file
sw172-22-46-174	deny	<input type="checkbox"/>	<input type="checkbox"/>	none	<input type="checkbox"/>	activeZoneSet	effectiveDB	file
sw172-22-46-222	deny	<input type="checkbox"/>	<input type="checkbox"/>	none	<input type="checkbox"/>	activeZoneSet	effectiveDB	file
sw172-22-46-153	deny	<input type="checkbox"/>	<input type="checkbox"/>	none	<input type="checkbox"/>	activeZoneSet	effectiveDB	file

- Step 3** In the Propagation column, choose **Full Zoneset** from the drop-down list.
- Step 4** Click **Apply Changes** to propagate the full zone set.
-

Enabling a One-Time Distribution

You can perform a one-time distribution of inactive, unmodified zone sets throughout the fabric.

To propagate a one-time distribution of the full zone set using Fabric Manager, perform this task:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Edit Local Full Zone Database dialog box.
- Step 2** Click the appropriate zone from the list in the left pane.
- Step 3** Click **Distribute** to distribute the full zone set across the fabric.
-

This procedure only distributes the full zone set information, it does not save the information to the startup configuration. To save the full zone set, you must explicitly save the running configuration to the startup configuration.



Note

The one-time distribution of the full zone set is supported in interop 2 and interop 3 modes, and not in interop 1 mode.

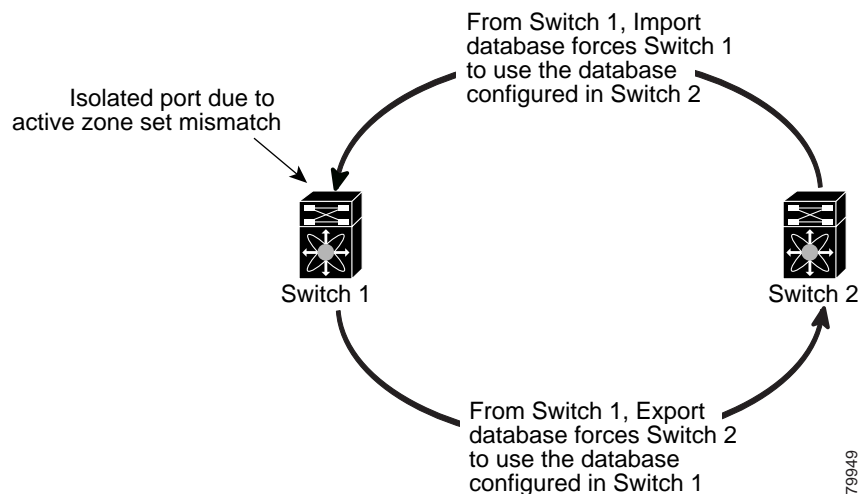
[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

About Recovering from Link Isolation

When two switches in a fabric are merged using a TE or E port, these TE and E ports may become isolated when the active zone set databases are different between the two switches or fabrics. When a TE port or an E port become isolated, you can recover that port from its isolated state using one of three options:

- Import the neighboring switch's active zone set database and replace the current active zone set (see [Figure 16-24](#)).
- Export the current database to the neighboring switch.
- Manually resolve the conflict by editing the full zone set, activating the corrected zone set, and then bringing up the link.

Figure 16-24 Importing and Exporting the Database



Importing and Exporting Zone Sets

To import or export the zone set information from or to an adjacent switch using Fabric Manager, perform this task:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Edit Local Full Zone Database dialog box.
- Step 2** Choose **Tools > Zone Merge Fail Recovery**.
You see the Zone Merge Failure Recovery dialog box as shown in [Figure 16-25](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 16-25 Zone Merge Failure Recovery Dialog Box



- Step 3** Click the **Import Active Zoneset** or the **Export Active Zoneset** radio button.
- Step 4** Choose the switch from which to import or export the zone set information from the drop-down list.
- Step 5** Choose the VSAN from which to import or export the zone set information from the drop-down list.
- Step 6** Choose the interface to use for the import process.
- Step 7** Click **OK** to import or export the active zone set.



Note

Enter the **import** and **export** from a single switch. Importing from one switch and exporting from another switch can lead to isolation again.

Zone Set Duplication

You can make a copy and then edit it without altering the existing active zone set. You can copy an active zone set from the bootflash: directory, volatile: directory, or slot0 to one of the following areas:

- To the full zone set
- To a remote location (using FTP, SCP, SFTP, or TFTP)

The active zone set is not part of the full zone set. You cannot make changes to an existing zone set and activate it if the full zone set is lost or is not propagated.



Caution

Copying an active zone set to a full zone set may overwrite a zone with the same name if it already exists in the full zone set database.

This section includes the following topics:

- [Copying Zone Sets, page 16-25](#)
- [About Backing Up and Restoring Zones, page 16-25](#)
- [Backing Up and Restoring Zones, page 16-25](#)
- [Renaming Zones, Zone Sets, and Aliases, page 16-26](#)
- [Cloning Zones, Zone Sets, FC Aliases, and Zone Attribute Groups, page 16-27](#)

Send comments to nx5000-docfeedback@cisco.com

- [Migrating a Non-MDS Database, page 16-28](#)
- [Clearing the Zone Server Database, page 16-28](#)

Copying Zone Sets


On Cisco Nexus 5000 Series switches, you cannot edit an active zone set. However, you can copy an active zone set to create a new zone set that you can edit.

To make a copy of a zone set using Fabric Manager, perform this task:

Step 1 Choose **Zone > Copy Full Zone Database**.

You see the Copy Full Zone Database dialog box as shown in [Figure 16-26](#).

Figure 16-26 Copy Full Zone Database Dialog Box



- Step 2** Click the **Active** or the **Full** radio button, depending on which type of database you want to copy.
- Step 3** Choose the source VSAN 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.
-

About Backing Up and Restoring Zones

You can back up the zone configuration to a workstation using TFTP. This zone backup file can then be used to restore the zone configuration on a switch. Restoring the zone configuration overwrites any existing zone configuration on a switch.

Backing Up and Restoring Zones

To back up or restore the full zone configuration using Fabric Manager, perform this task:

Step 1 Choose **Zone > Edit Local Full Zone Database**.

Send comments to nx5000-docfeedback@cisco.com

You see the Select VSAN dialog box.


Step 2 Choose a VSAN and click **OK**.

You see the Edit Local Full Zone Database dialog box for the selected VSAN.

Step 3 Choose **File > Backup** to back up the existing zone configuration to a workstation using TFTP, or choose **File > Restore** to restore a saved zone configuration.

You see the Restore Zone Configuration dialog box as shown in [Figure 16-27](#).

Figure 16-27 Restore Zone Configuration Dialog Box



You can edit this configuration before restoring it to the switch.

Step 4 Click **Continue Restore**, or click **Close** to close the dialog box without restoring.

Renaming Zones, Zone Sets, and Aliases

To rename a zone, zone set, or alias using Fabric Manager, perform this task:

Step 1 Choose **Zone > Edit Local Full Zone Database**.

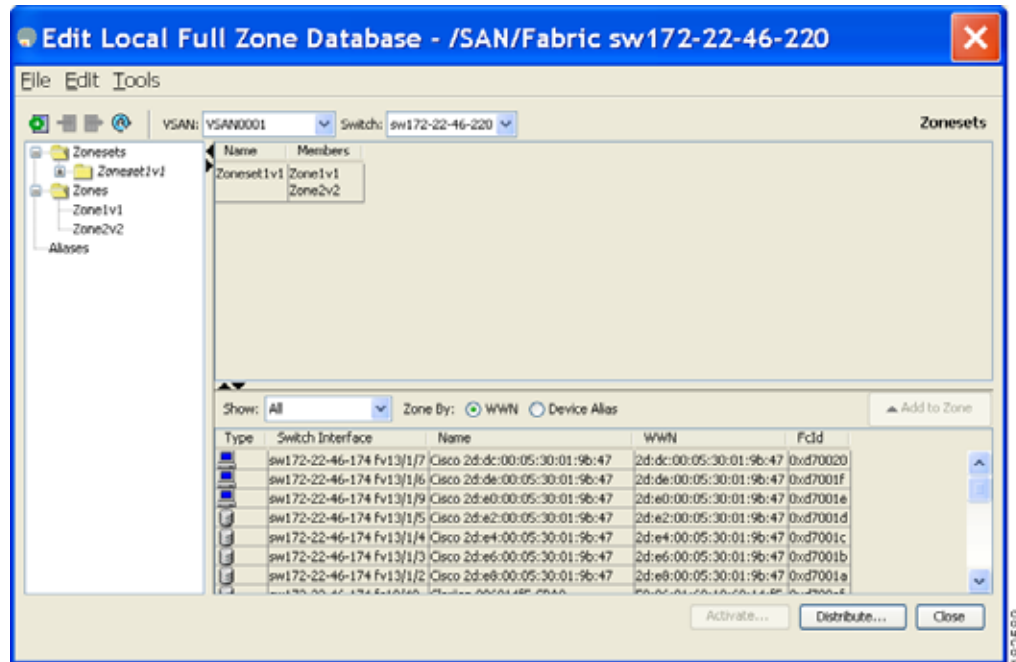
You see the Select VSAN dialog box.

Step 2 Choose a VSAN and click **OK**.

You see the Edit Local Full Zone Database dialog box for the selected VSAN as shown in [Figure 16-28](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 16-28 Edit Local Full Zone Database Dialog Box



- Step 3** Click a zone or zone set in the left pane.
- Step 4** Choose **Edit > Rename**.
An edit box appears around the zone or zone set name.
- Step 5** Enter a new name.
- Step 6** Click **Activate** or **Distribute**.


Cloning Zones, Zone Sets, FC Aliases, and Zone Attribute Groups

To clone a zone, zone set, fcalias, or zone-attribute-group, perform this task:

- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Choose a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
- Step 3** Choose **Edit > Clone**.
You see the Clone Zoneset dialog box as shown in [Figure 16-29](#). The default name is the word **Clone** followed by the original name.

Send comments to nx5000-docfeedback@cisco.com

Figure 16-29 Clone Zoneset Dialog Box



Step 4 Change the name for the cloned entry.

Step 5 Click **OK** to save the new clone.

The cloned database now appears along with the original database.

Migrating a Non-MDS Database

To use the Zone Migration Wizard to migrate a non-MDS database using Fabric Manager, perform this task:

Step 1 Choose **Zone > Migrate Non-MDS Database**.

You see the Zone Migration Wizard.

Step 2 Follow the prompts in the wizard to migrate the database.

Clearing the Zone Server Database

You can clear all configured information in the zone server database for the specified VSAN.

To clear the zone server database, see the *Cisco Cisco Nexus 5000 Series CLI Configuration Guide*.



Note

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



Note

After clearing the zone server database, you must explicitly copy the running configuration to the startup configuration to save the configuration.

Verifying Zone Information

To view zone information and statistics using Fabric Manager, perform this task:

Step 1 Expand a **VSAN** and click a zone set in the Logical Domains pane.

You see the zone configuration in the Information pane.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

- Step 2** Click the **Read Only Violations, Statistics** tab or the **LUN Zoning Statistics** tab to view statistics for the selected zone.

Enhanced Zoning

The zoning feature complies with the FC-GS-4 and FC-SW-3 standards. Both standards support the basic zoning functionalities explained in the previous section and the enhanced zoning functionalities described in this section.

This section includes the following topics:

- [About Enhanced Zoning, page 16-29](#)
- [Changing from Basic Zoning to Enhanced Zoning, page 16-30](#)
- [Changing from Enhanced Zoning to Basic Zoning, page 16-30](#)
- [Enabling Enhanced Zoning, page 16-31](#)
- [Merging the Database, page 16-31](#)
- [Analyzing a Zone Merge, page 16-32](#)
- [Configuring Zone Merge Control Policies, page 16-32](#)

About Enhanced Zoning

[Table 16-2](#) lists the advantages of the enhanced zoning feature in all switches in the Cisco Nexus 5000 Series.

Table 16-2 Advantages of Enhanced Zoning

Basic Zoning	Enhanced Zoning	Enhanced Zoning Advantages
Administrators can make simultaneous configuration changes. Upon activation, one administrator can overwrite another administrator's changes.	Performs all configurations within a single configuration session. When you begin a session, the switch locks the entire fabric to implement the change.	One configuration session for the entire fabric to ensure consistency within the fabric.
If a zone is part of multiple zone sets, you create an instance of this zone in each zone set	References to the zone are used by the zone sets as required once you define the zone.	Reduced payload size as the zone is referenced. The size is more pronounced with bigger databases.
The default zone policy is defined per switch. To ensure smooth fabric operation, all switches in the fabric must have the same default zone setting.	Enforces and exchanges the default zone setting throughout the fabric.	Fabric-wide policy enforcement reduces troubleshooting time.
To retrieve the results of the activation on a per switch basis, the managing switch provides a combined status about the activation. It does not identify the failure switch.	Retrieves the activation results and the nature of the problem from each remote switch.	Enhanced error reporting eases the troubleshooting process

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Table 16-2 Advantages of Enhanced Zoning (continued)

Basic Zoning	Enhanced Zoning	Enhanced Zoning Advantages
To distribute the zoning database, you must reactivate the same zone set. The reactivation may affect hardware changes for hard zoning on the local switch and on remote switches.	Implements changes to the zoning database and distributes it without reactivation.	Distribution of zone sets without activation avoids hardware changes for hard zoning in the switches.
The Cisco-specific zone member types (symbolic node name, and other types) may be used by other non-Cisco switches. During a merge, the Cisco-specific types can be misunderstood by the non-Cisco switches.	Provides a vendor ID along with a vendor-specific type value to uniquely identify a member type.	Unique vendor type.
The fWWN-based zone membership is only supported in Cisco interop mode.	Supports fWWN-based membership in the standard interop mode (interop mode 1).	The fWWN-based member type is standardized.

Changing from Basic Zoning to Enhanced Zoning

To change to the enhanced zoning mode from the basic mode, perform this task:

-
- Step 1** Verify that all switches in the fabric are capable of working in the enhanced mode.
- If one or more switches are not capable of working in enhanced mode, then your request to move to enhanced mode is rejected.
- Step 2** Set the operation mode to enhanced zoning mode.
- You will automatically start a session, acquire a fabric wide lock, distribute the active and full zoning database using the enhanced zoning data structures, distribute zoning policies and then release the lock. All switches in the fabric then move to the enhanced zoning mode.



Tip After moving from basic zoning to enhanced zoning, we recommend that you save the running configuration.

Changing from Enhanced Zoning to Basic Zoning

Cisco SAN switches allow you to change from enhanced zoning to basic zoning to enable you to downgrade and upgrade to other Cisco NX-OS releases.

To change to the basic zoning mode from the enhanced mode, perform this task:

-
- Step 1** Verify that the active and full zone set do not contain any configuration that is specific to the enhanced zoning mode.
- If such configurations exist, delete them before proceeding with this procedure. If you do not delete the existing configuration, the switch software automatically removes them.
- Step 2** Set the operation mode to basic zoning mode.

Send comments to nx5000-docfeedback@cisco.com

You will automatically start a session, acquire a fabric-wide lock, distribute the zoning information using the basic zoning data structure, apply the configuration changes and release the lock from all switches in the fabric. All switches in the fabric then move to basic zoning mode.

Enabling Enhanced Zoning

By default, the enhanced zoning feature is disabled in all switches in the Cisco Nexus 5000 Series.

To enable enhanced zoning in a VSAN using Fabric Manager, perform this task:

-
- Step 1** In the Logical Domains pane, expand a VSAN, and then choose a zone set.
You see the zone set configuration in the Information pane.
- Step 2** Click the **Enhanced** tab.
You see the current enhanced zoning configuration.
- Step 3** In the Action drop-down list, choose **enhanced** to enable enhanced zoning in this VSAN.
- Step 4** Click **Apply Changes** to save these changes.
-

Merging the Database

The merge method depends on the fabric-wide merge control setting:

- **Restrict**—If the two databases are not identical, the ISLs between the switches are isolated.
- **Allow**—The two databases are merged using the merge rules specified in [Table 16-3](#).

Table 16-3 Database Zone Merge Status

Local Database	Adjacent Database	Merge Status	Results of the Merge
The databases contain zone sets with the same name ¹ but different zones, aliases, and attributes groups.		Successful.	The union of the local and adjacent databases.
The databases contains a zone, zone alias, or zone attribute group object with same name ¹ but different members.		Failed.	ISLs are isolated.
Empty.	Contains data.	Successful.	The adjacent database information populates the local database.
Contains data.	Empty.	Successful.	The local database information populates the adjacent database.

1. In the enhanced zoning mode, the active zone set does not have a name in interop mode 1. The zone set names are only present for full zone sets.

Send comments to nx5000-docfeedback@cisco.com

The merge process operates as follows:

1. The software compares the protocol versions. If the protocol versions differ, then the ISL is isolated.
2. If the protocol versions are the same, then the zone policies are compared. If the zone policies differ, then the ISL is isolated.
3. If the zone merge options are the same, then the comparison is implemented based on the merge control setting.
 - a. If the setting is restrict, the active zone set and the full zone set should be identical. Otherwise, the link is isolated.
 - b. If the setting is allow, then the merge rules are used to perform the merge.

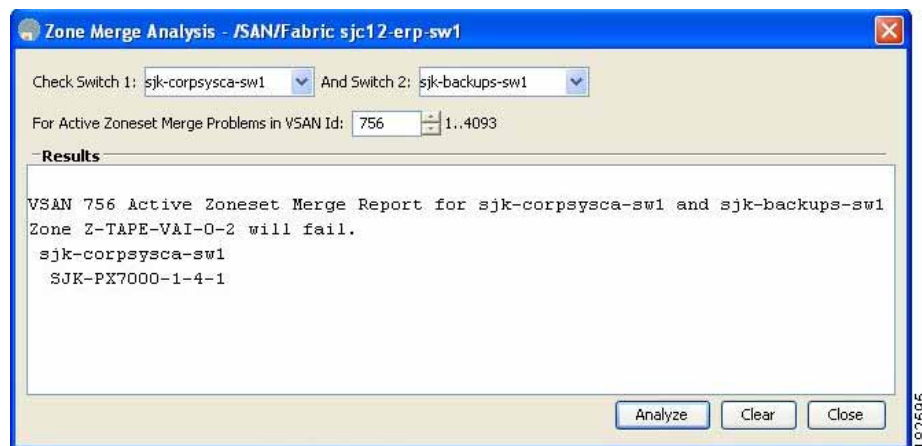
Analyzing a Zone Merge

To perform a zone merge analysis using Fabric Manager, perform this task:

-
- Step 1** Choose **Zone > Merge Analysis**.

You see the Zone Merge Analysis dialog box as shown in [Figure 16-30](#).

Figure 16-30 Zone Merge Analysis Dialog Box



- 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 Id field.
- Step 5** Click **Analyze** to analyze the zone merge.
- Step 6** Click **Clear** to clear the analysis data in the Zone Merge Analysis dialog box.
-

Configuring Zone Merge Control Policies

To configure merge control policies, see the *Cisco Cisco Nexus 5000 Series CLI Configuration Guide*.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Compacting the Zone Database

You can delete excess zones and compact the zone database for the VSAN.



Note

A merge failure occurs when a switch supports more than 2000 zones per VSAN but its neighbor does not. Also, zone set activation can fail if the switch has more than 2000 zones per VSAN and not all switches in the fabric support more than 2000 zones per VSAN.

To compact the zone database for downgrading, see the *Cisco Cisco Nexus 5000 Series CLI Configuration Guide*.

Default Settings

Table 16-4 lists the default settings for basic zone parameters.

Table 16-4 Default Basic Zone Parameters

Parameters	Default
Default zone policy	Denied to all members.
Full zone set distribute	The full zone set(s) is not distributed.
Enhanced zoning	Disabled.

Send comments to nx5000-docfeedback@cisco.com



Distributing Device Alias Services

Switches in the Cisco Nexus 5000 Series support Distributed Device Alias Services (device aliases) on a fabric-wide basis.

This chapter includes the following sections:

- [Information About Device Aliases, page 17-1](#)
- [Device Alias Databases, page 17-3](#)
- [Legacy Zone Alias Conversion, page 17-7](#)
- [Database Merge Guidelines, page 17-8](#)
- [Default Settings, page 17-9](#)

Information About Device Aliases

When the port WWN (pWWN) of a device must be specified to configure features (for example, zoning, DPVM, or port security) in a Cisco Nexus 5000 Series switch, you must assign the correct device name each time you configure these features. An inaccurate device name may cause unexpected results. You can circumvent this problem if you define a user-friendly name for a pWWN and use this name in all the configuration commands as required. These user-friendly names are referred to as *device aliases* in this chapter.

This section includes the following topics:

- [Device Alias Features, page 17-1](#)
- [Device Alias Requirements, page 17-2](#)
- [Zone Aliases Versus Device Aliases, page 17-2](#)

Device Alias Features

Device aliases have the following features:

- The device alias information is independent of the VSAN configuration.
- The device alias configuration and distribution is independent of the zone server and the zone server database.
- You can import legacy zone alias configurations without losing data.

Send comments to nx5000-docfeedback@cisco.com

- The device alias application uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management and distribution. Device aliases use the coordinated distribution mode and the fabric-wide distribution scope (see [Chapter 7, “Using Cisco Fabric Services”](#)).
- Basic and enhanced modes. See the [“Device Alias Modes” section on page 17-3](#).
- Device aliases used to configure zones, IVR zones, or port security features are displayed automatically with their respective pWWNs in the **show** command output.

Device Alias Requirements

Device aliases have the following requirements:

- You can only assign device aliases to pWWNs.
- There must be a one-to-one relationship between the pWWN and the device alias that maps to it.
- A device alias name is restricted to 64 alphanumeric characters and may include one or more of the following characters:
 - a to z and A to Z
 - Device alias names must begin with an alphabetic character (a to z or A to Z).
 - 1 to 9
 - - (hyphen) and _ (underscore)
 - \$ (dollar sign) and ^ (up caret)

Zone Aliases Versus Device Aliases

[Table 17-1](#) compares the configuration differences between zone-based alias configuration and device alias configuration.

Table 17-1 Comparison Between Zone Aliases and Device Aliases

Zone-Based Aliases	Device Aliases
Aliases are limited to the specified VSAN.	You can define device aliases without specifying the VSAN number. You can also use the same definition in one or more VSANs without any restrictions.
Zone aliases are part of the zoning configuration. The alias mapping cannot be used to configure other features.	Device aliases can be used with any feature that uses the pWWN.
You can use any zone member type to specify the end devices.	Only pWWNs are supported.
Configuration is contained within the zone server database and is not available to other features.	Device aliases are not restricted to zoning. Device alias configuration is available to the FCNS, zone, fcping, and traceroute applications.

Send comments to nx5000-docfeedback@cisco.com

Device Alias Databases

The device alias feature uses two databases to accept and implement device alias configurations.

- Effective database—The database currently used by the fabric.
- Pending database—Your subsequent device alias configuration changes are stored in the pending database.

If you modify the device alias configuration, you need to commit or discard the changes as the fabric remains locked during this period.

This section includes the following topics:

- [Device Alias Modes, page 17-3](#)
- [Changing Device Alias Mode Guidelines, page 17-3](#)
- [Configuring Device Alias Modes, page 17-4](#)
- [About Device Alias Distribution, page 17-5](#)
- [Distributing the Device Alias Database, page 17-5](#)
- [About Creating a Device Alias, page 17-5](#)
- [Creating a Device Alias, page 17-6](#)
- [Committing Changes, page 17-6](#)
- [Discarding Changes, page 17-7](#)

Device Alias Modes

You can specify that aliases operate in basic or enhanced modes.

When operating in basic mode, which is the default mode, the device alias is immediately expanded to a pWWN. In basic mode, when device aliases are changed to point to a new HBA, for example, that change is not reflected in the zone server. Users must remove the previous HBA's pWWN, add the new HBA's pWWN, and then reactivate the zoneset.

When operating in enhanced mode, applications accept a device alias name in its “native” format. Instead of expanding the device alias to a pWWN, the device alias name is stored in the configuration and distributed in its native device alias format. So applications such as zone server, PSM or DPVM can automatically keep track of the device alias membership changes and enforce them accordingly. The primary benefit of operating in enhanced mode is that you have a single point of change.

Whenever you change device alias modes, the change is distributed to other switches in the network only if device alias distribution is enabled or on. Otherwise, the mode change only takes place on the local switch.



Note

Enhanced mode, or native device alias-based configurations are not accepted in interop mode VSANs. IVR zoneset activation will fail in interop mode VSANs if the corresponding zones have native device alias-based members.

Changing Device Alias Mode Guidelines

When changing device alias modes, follow these guidelines:

Send comments to nx5000-docfeedback@cisco.com

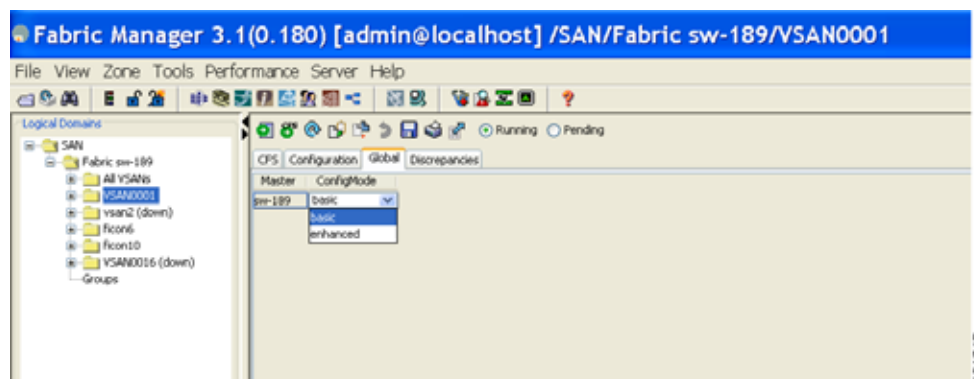
- If two fabrics running in different device alias modes are joined together, the device alias merge will fail. There is no automatic conversion to one mode or the other during the merge process. In this situation, you must select one mode over the other.
- Before changing from enhanced to basic mode, you must first explicitly remove all native device alias-based configurations from both local and remote switches, or, replace all device alias-based configuration members with the corresponding pWWN.
- If you remove a device alias from the device alias database, all applications will automatically stop enforcing the corresponding device alias. If that corresponding device alias is part of an active zoneset, all the traffic to and from that pWWN is disrupted.
- Renaming the device alias not only changes the device alias name in the device alias database, but also replaces the corresponding device alias configuration in all the applications.
- When a new device alias is added to the device alias database, and the application configuration is present on that device alias, it automatically takes effect. For example, if the corresponding device alias is part of the active zoneset and the device is online, then zoning is enforced automatically. You do not have to reactivate the zoneset.
- If a device alias name is mapped to a new HBA's pWWN, then the application's enforcement changes accordingly. In this case, the zone server automatically enforces zoning based on the new HBA's pWWN.

Configuring Device Alias Modes

To configure device aliases to operate in enhanced mode using Fabric Manager, perform this task:

-
- Step 1** Expand **End Devices**, and then choose **Device Alias** in the Physical Attributes pane. You see the device alias configuration in the Information pane as shown in [Figure 17-1](#).
 - Step 2** Click the **Mode** tab.
 - Step 3** Choose **enhanced** from the ConfigMode drop-down list.

Figure 17-1 Configuring Modes



- Step 4** Click **Apply Changes** to commit and distribute these changes, or click **Undo Changes** to discard any unsaved changes.
-

Send comments to nx5000-docfeedback@cisco.com

About Device Alias Distribution

By default, device alias distribution is enabled. The device alias feature uses CFS to distribute the modifications to all switches in a fabric.

If device alias distribution is disabled, database changes are not distributed to the switches in the fabric. The same changes would have to be performed manually on all switches in the fabric to keep the device alias database up-to-date. Database changes immediately take effect, so there would not be any pending database and commit or abort operations either. If you have not committed the changes and you disable distribution, then a commit task will fail.

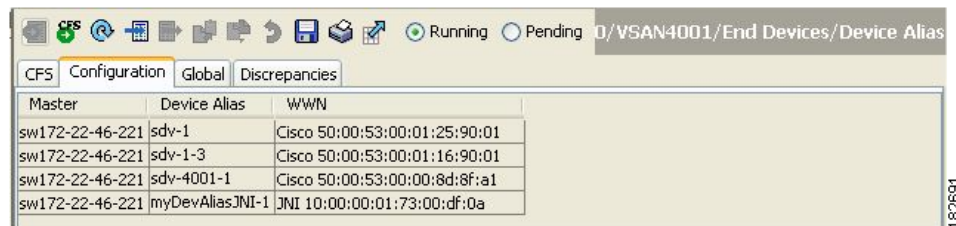
The following example displays a failed device alias status:

Distributing the Device Alias Database

To enable the device alias distribution using Fabric Manager, perform this task:

- Step 1** Expand **End Devices**, and then choose **Device Alias** in the Physical Attributes pane. You see the device alias configuration in the Information pane as shown in [Figure 17-2](#).

Figure 17-2 Device Aliases in Fabric Manager



Master	Device Alias	WWN
sw172-22-46-221	sdv-1	Cisco 50:00:53:00:01:25:90:01
sw172-22-46-221	sdv-1-3	Cisco 50:00:53:00:01:16:90:01
sw172-22-46-221	sdv-4001-1	Cisco 50:00:53:00:00:8d:8f:a1
sw172-22-46-221	myDevAliasJMI-1	JMI 10:00:00:01:73:00:df:0a

- Step 2** Click the **CFS** tab.
- Step 3** Choose **enable** from the Feature Admin column to enable switch aliases.
- Step 4** Choose **commitChanges** from the Config Action column for the newly enabled switches.
- Step 5** Click **Apply Changes** to commit and distribute these changes, or click **Undo Changes** to discard any unsaved changes.

About Creating a Device Alias

When you perform any device alias configuration task (regardless of which device alias task), the fabric is automatically locked for the device alias feature. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the effective database is obtained and used as the pending database. Subsequent modifications are made to the pending database. The pending database remains in use until you commit the modifications to the pending database or discard (**abort**) the changes to the pending database.

Send comments to nx5000-docfeedback@cisco.com

Creating a Device Alias

To lock the fabric using Fabric Manager, perform this task:

**Note**

You create a device alias for a locked fabric in the pending database.


Step 1 Expand **End Devices**, and then choose **Device Alias** in the Physical Attributes pane.

You see the device alias configuration in the Information pane.

Step 2 Click the **Configuration** tab and click the **Create Row** icon.

You see the Create Device Alias dialog box as shown in [Figure 17-3](#).

Figure 17-3 Create Device Alias Dialog Box



Step 3 Complete the Alias name and pWWN fields.

Step 4 Click **Create** to create this alias.

Committing Changes

If you commit the changes made to the pending database, the following events occur:

1. The pending database content overwrites the effective database content.
2. The pending database is distributed to the switches in the fabric and the effective database on those switches is overwritten with the new changes.
3. The pending database is emptied of its contents.
4. The fabric lock is released for this feature.

To commit the changes to the device alias database using Fabric Manager, perform this task:

Step 1 Expand **End Devices**, and then choose **Device Alias** in the Physical Attributes pane.

You see the device alias configuration in the Information pane.

Step 2 Click the **CFS** tab.

Step 3 Choose **enable** from the Feature Admin column to enable switch aliases.

Step 4 Choose **commitChanges** from the Config Action column for the newly enabled switches.

Send comments to nx5000-docfeedback@cisco.com

- Step 5** Click **Apply Changes** to commit and distribute these changes, or click **Undo Changes** to discard any unsaved changes.
-

Discarding Changes

If you discard the changes made to the pending database, the following events occur:

1. The effective database contents remain unaffected.
2. The pending database is emptied of its contents.
3. The fabric lock is released for this feature.

To discard the device alias session using Fabric Manager, perform this task:

- Step 1** Expand **End Devices**, and then choose **Device Alias** in the Physical Attributes pane.
You see the device alias configuration in the Information pane.
- Step 2** Click the **CFS** tab.
- Step 3** Choose **abort** from the Config Action column for the newly enabled switches.
- Step 4** Click **Apply Changes** to discard the changes.
-

Legacy Zone Alias Conversion

You can import legacy zone alias configurations to use this feature without losing data if they satisfy the following restrictions:

- Each zone alias has only one member.
- The member type is pWWN.

If any name or definition conflict exists, the zone aliases are not imported.



Tip

Ensure that you copy any required zone aliases to the device alias database as required by your configuration.

When an import operation is complete, the modified alias database is distributed to all other switches in the physical fabric when you perform the **commit** operation. If you do not want to distribute the configuration to other switches in the fabric, you can perform the **abort** operation and the merge changes are completely discarded.

Using Device Aliases or FC Aliases

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

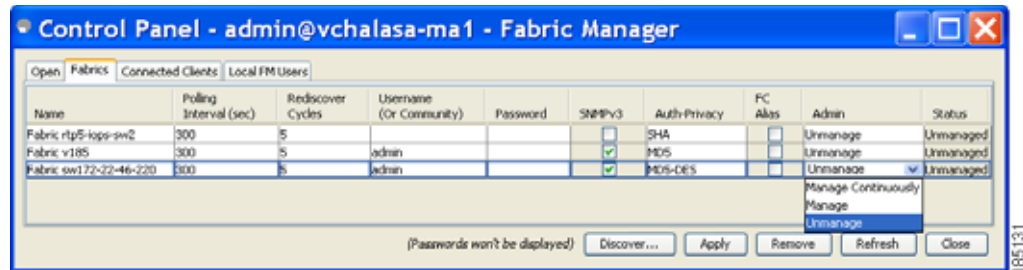
Send comments to nx5000-docfeedback@cisco.com

To change whether Fabric Manager uses FC aliases or global device aliases, perform this task:

Step 1 Click **Server > Admin**.

You see the Control Panel dialog box with the Fabrics tab open (see [Figure 17-4](#)).

Figure 17-4 Control Panel Dialog Box



Step 2 Check the **FC Alias** check box to use FC aliases or uncheck to use global device aliases for each fabric that you are managing with Fabric Manager Server.

Step 3 Click **Apply** to save these changes.

Database Merge Guidelines

When merging two device alias databases, follow these guidelines:

- Verify that two device aliases with different names are not mapped to the same pWWN.
- Verify that two identical pWWNs are not mapped to two different device aliases.
- Verify that the combined number of device aliases in both databases does not exceed 8K (8191 device aliases) in fabrics running Cisco MDS SAN-OS release 3.0 (x) and earlier, and 20K in fabrics running Cisco MDS SAN-OS release 3.1(x) and later.

If the combined number of device entries in both databases exceeds the supported configuration limit, then the merge will fail. For example, if database N has 6000 device aliases and database M has 2192 device aliases, and you are running SAN-OS 3.0(x) or earlier, then this merge operation will fail. Merge operations will also fail if there is a device alias mode mismatch.

For additional information, see the [“CFS Merge Support” section on page 7-6](#).

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Default Settings

Table 17-2 lists the default settings for device alias parameters.

Table 17-2 *Default Device Alias Parameters*

Parameters	Default
Device alias distribution	Enabled.
Device alias mode	Basic.
Database in use	Effective database.
Database to accept changes	Pending database.
Device alias fabric lock state	Locked with the first device alias task.

Send comments to nx5000-docfeedback@cisco.com



CHAPTER 18

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 the E mode and TE mode Fibre Channel interfaces on Cisco Nexus 5000 Series 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. FSPF provides the following capabilities:

- Dynamically computes routes throughout a fabric by establishing the shortest and quickest path between any two switches.
- Selects an alternative path in the event of the failure of a given path. FSPF supports multiple paths and automatically computes an alternative path around a failed link. It provides a preferred route when two equal paths are available.

This chapter provides details on Fibre Channel routing services and protocols. It includes the following sections:

- [Information About FSPF, page 18-1](#)
- [FSPF Global Configuration, page 18-3](#)
- [FSPF Interface Configuration, page 18-5](#)
- [FSPF Routes, page 18-11](#)
- [In-Order Delivery, page 18-12](#)
- [Default Settings, page 18-16](#)

Information About FSPF

FSPF is the protocol currently standardized by the T11 committee for routing in Fibre Channel networks. The FSPF protocol has the following characteristics and features:

- Supports multipath routing.
- Bases path status on a link state protocol.
- Routes hop by hop, based only on the domain ID.
- Runs only on E ports or TE ports and provides a loop free topology.
- Runs on a per VSAN basis. Connectivity in a given VSAN in a fabric is guaranteed only for the switches configured in that VSAN.

Send comments to nx5000-docfeedback@cisco.com

- Uses a topology database to keep track of the state of the links on all switches in the fabric and associates a cost with each link.
- Guarantees a fast reconvergence time in case of a topology change. Uses the standard Dijkstra algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.

FSPF Examples

This section provides examples of topologies and applications that demonstrate the benefits of FSPF.




Note

The FSPF feature can be used on any topology.

Fault Tolerant Fabric Example

Figure 18-1 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 18-1 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 Link Example

To improve on the topology in Figure 18-1, each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. Figure 18-2 shows this arrangement. Because switches in the Cisco Nexus 5000 Series support port channels, 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 SAN port channel. This configuration also improves the resiliency of the network. The failure of a link in a SAN port channel does not trigger a route change, which reduces the risks of routing loops, traffic loss, or fabric downtime for route reconfiguration.

Send comments to nx5000-docfeedback@cisco.com

Figure 18-2 Fault Tolerant Fabric with Redundant Links



For example, if all links are of equal speed and no SAN port channels 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 SAN port channels exist, these paths are reduced to two.

FSPF Global Configuration

By default, FSPF is enabled on switches in the Cisco Nexus 5000 Series.

Some FSPF features can be globally configured in each VSAN. By configuring a feature for the entire VSAN, you do not have to specify the VSAN number for every command. This global configuration feature also reduces the chance of typing errors or other minor configuration errors.



Note

FSPF is enabled by default. Generally, you do not need to configure these advanced features.



Caution

The default for the backbone region is 0 (zero). You do not need to change this setting unless your region is different from the default. If you are operating with other vendors using the backbone region, you can change this default to be compatible with those settings.

This section includes the following topics:

- [About SPF Computational Hold Times, page 18-3](#)
- [About Link State Records, page 18-3](#)
- [Configuring FSPF on a VSAN, page 18-4](#)
- [Resetting FSPF to the Default Configuration, page 18-5](#)
- [Enabling or Disabling FSPF, page 18-5](#)

About SPF Computational Hold Times

The SPF computational hold time sets the minimum time between two consecutive SPF computations on the VSAN. Setting this to a small value means that FSPF reacts faster to any fabric changes by recomputing paths on the VSAN. A small SPF computational hold time uses more switch CPU time.

About Link State Records

Each time a new switch enters the fabric, a link state record (LSR) is sent to the neighboring switches, and then flooded throughout the fabric.

Send comments to nx5000-docfeedback@cisco.com

Table 18-1 displays the default settings for switch responses.

Table 18-1 LSR Default Settings

LSR Option	Default	Description
Acknowledgment interval (RxmtInterval)	5 seconds	The time a switch waits for an acknowledgment from the LSR before retransmission.
Refresh time (LSRefreshTime)	30 minutes	The time a switch waits before sending an LSR refresh transmission.
Maximum age (MaxAge)	60 minutes	The time a switch waits before dropping the LSR from the database.

The LSR minimum arrival time is the period between receiving LSR updates on this VSAN. Any LSR updates that arrive before the LSR minimum arrival time are discarded.

The LSR minimum interval time is the frequency at which this switch sends LSR updates on a VSAN.

Configuring FSPF on a VSAN

To configure an FSPF feature for the entire VSAN using Fabric Manager, perform this task:

- Step 1** Expand a Fabric, expand a VSAN, and then choose **FSPF** for a VSAN that you want to configure for FSPF.

You see the FSPF configuration in the Information pane as shown in [Figure 18-3](#).

Figure 18-3 FSPF General Information

Switch	Status Admin	Status Oper	Set To Default	RegionID	DomainID	Spf Comp. HoldTime	Spf Comp. Delay	LSR Min Arrival(ms)	LSR Min Interval(ms)	LSR Refresh Time(min)	LSR Max Age(min)	CreateTime
sw172-22-46-223	up	up	<input type="checkbox"/>	0	0vea(236)	0	0	1000	2000	30	60	2007/03/29-1
sw172-22-46-224	up	up	<input type="checkbox"/>	0	0vea(234)	0	0	1000	2000	30	60	2007/03/14-0
sw172-22-46-220	up	up	<input type="checkbox"/>	0	0vea(239)	0	0	1000	2000	30	60	2007/04/04-1
sw172-22-46-221	up	up	<input type="checkbox"/>	0	0vea(238)	0	0	1000	2000	30	60	2007/03/27-1
sw172-22-46-222	up	up	<input type="checkbox"/>	0	0vea(233)	0	0	1000	2000	30	60	2007/03/14-0
sw172-22-46-223	up	up	<input type="checkbox"/>	0	0vea(235)	0	0	1000	2000	30	60	2007/03/14-0
sw172-22-46-225	up	up	<input type="checkbox"/>	0	0vea(232)	0	0	1000	2000	30	60	2007/03/29-1
sw172-22-46-174	up	up	<input type="checkbox"/>	0	0vea(237)	0	0	1000	2000	30	60	2007/03/14-0

You can enter values for RegionID, Spf Comp Holdtime, LSR Min Arrival, and LSR Min Interval. These values apply to all interfaces on the VSAN.

- Step 2** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

Send comments to nx5000-docfeedback@cisco.com

Resetting FSPF to the Default Configuration

To return the FSPF VSAN global configuration to its factory default using Fabric Manager, perform this task:

-
- Step 1** Expand a Fabric, expand a VSAN and then choose **FSPF** for a VSAN that you want to configure for FSPF.
- You see the FSPF configuration in the Information pane as shown in [Figure 18-3](#).
- Step 2** Check the **SetToDefault** check box for a switch.
- Step 3** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-

Enabling or Disabling FSPF

To enable or disable FSPF using Fabric Manager, perform this task:

-
- Step 1** Expand a Fabric, expand a VSAN, and then choose **FSPF** for a VSAN that you want to configure for FSPF.
- You see the FSPF configuration in the Information pane as shown in [Figure 18-3](#).
- Step 2** Set the Status Admin drop-down list to **up** to enable FSPF or to **down** to disable FSPF.
- Step 3** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-

FSPF Interface Configuration

Several FSPF commands are available on a per-interface basis. These configuration procedures apply to an interface in a specific VSAN.

This section includes the following topics:

- [About FSPF Link Cost, page 18-6](#)
- [Configuring FSPF Link Cost, page 18-6](#)
- [About Hello Time Intervals, page 18-6](#)
- [Configuring Hello Time Intervals, page 18-7](#)
- [About Dead Time Intervals, page 18-7](#)
- [Configuring Dead Time Intervals, page 18-7](#)
- [About Retransmitting Intervals, page 18-7](#)
- [Configuring Retransmitting Intervals, page 18-8](#)
- [About Disabling FSPF for Specific Interfaces, page 18-8](#)
- [Disabling FSPF for Specific Interfaces, page 18-8](#)
- [Displaying the FSPF Database, page 18-9](#)

Send comments to nx5000-docfeedback@cisco.com

- [Viewing FSPF Statistics, page 18-10](#)

About FSPF Link Cost

FSPF tracks the state of links on all switches in the fabric, associates a cost with each link in its database, and then chooses the path with a minimal cost. The cost associated with an interface can be administratively changed to implement the FSPF route selection. The integer value to specify cost can range from 1 to 65,535. The default cost for 1 Gbps is 1000 and for 2 Gbps is 500.

Configuring FSPF Link Cost

To configure FSPF link cost using Fabric Manager, perform this task:

- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**. You see the interface configuration in the Information pane.
- Step 2** Click the **FSPF** tab. You see the FSPF interface configuration in the Information pane as shown in [Figure 18-4](#).

Figure 18-4 Fibre Channel Physical FSPF Interface

Switch	VSAN id	Interface	Set To Default	Cost	Admin Status	Hello Interval	Dead Interval	Retx Interval	Neighbor State	Neighbor Domain	Neighbor PortIndex	CreateTime
sw172-22-46-182	1	fc1/16	<input type="checkbox"/>	500	up	20	80	5	full	0x0a(218)	0x10001	2006/03/10-15:44:24
sw172-22-46-224	1	fc1/5	<input type="checkbox"/>	500	up	20	80	5	full	0x07(215)	0x10004	2006/03/12-20:24:30
sw172-22-46-220	1	fc1/1	<input type="checkbox"/>	250	up	20	80	5	full	0x02(210)	0x10300	2006/03/12-20:19:46
sw172-22-46-225	1	fc1/5	<input type="checkbox"/>	500	up	20	80	5	full	0x09(217)	0x10004	2006/03/12-20:24:42
sw172-22-46-224	1	fc1/9	<input type="checkbox"/>	500	up	20	80	5	full	0x07(215)	0x10008	2006/03/12-20:24:48
sw172-22-46-225	1	fc1/9	<input type="checkbox"/>	500	up	20	80	5	full	0x08(217)	0x10008	2006/03/12-20:24:42
sw172-22-46-220	1	fc1/2	<input type="checkbox"/>	500	up	20	80	5	full	0x02(210)	0x10306	2006/03/12-20:19:46
sw172-22-46-224	1	fc1/3	<input type="checkbox"/>	500	up	20	80	5	full	0x07(215)	0x1000c	2006/03/12-20:24:48
sw172-22-46-225	1	fc1/3	<input type="checkbox"/>	500	up	20	80	5	full	0x09(217)	0x1000c	2006/03/12-20:24:42
sw172-22-46-220	1	fc1/3	<input type="checkbox"/>	250	up	20	80	5	full	0x08(219)	0x1090c	2006/03/10-15:45:01
sw172-22-46-224	1	fc1/21	<input type="checkbox"/>	500	up	20	80	5	full	0x08(218)	0x10008	2006/03/10-15:45:01
sw172-22-46-225	4001	fc1/5	<input type="checkbox"/>	500	up	20	80	5	full	0x0b(235)	0x10004	2006/03/12-20:24:43
sw172-22-46-220	1	fc1/4	<input type="checkbox"/>	250	up	20	80	5	full	0x08(219)	0x10904	2006/03/12-21:06:00
sw172-22-46-153	1	fc1/9	<input type="checkbox"/>	500	up	20	80	5	full	0x08(217)	0x10014	2006/03/10-15:45:01

- Step 3** Double-click in the Cost field of a switch and change the value.
- Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

About Hello Time Intervals

You can set the FSPF Hello time interval to specify the interval between the periodic hello messages sent to verify the health of the link. The integer value can range from 1 to 65,535 seconds.



Note

This value must be the same in the ports at both ends of the ISL.

Send comments to nx5000-docfeedback@cisco.com

Configuring Hello Time Intervals

To configure the FSPF Hello time interval using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**.
You see the interface configuration in the Information pane.
 - Step 2** Click the **FSPF** tab.
You see the FSPF interface configuration in the Information pane as shown in [Figure 18-4](#).
 - Step 3** Change the Hello Interval field for a switch.
 - Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-

About Dead Time Intervals

You can set the FSPF dead time interval to specify the maximum interval for which a hello message must be received before the neighbor is considered lost and removed from the database. The integer value can range from 1 to 65,535 seconds.



Note

This value must be the same in the ports at both ends of the ISL.

Configuring Dead Time Intervals

To configure the FSPF dead time interval using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**.
You see the interface configuration in the Information pane.
 - Step 2** Click the **FSPF** tab.
You see the FSPF interface configuration in the Information pane as shown in [Figure 18-4](#).
 - Step 3** Double-click the Dead Interval field for a switch and provide a new value.
 - Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-

About Retransmitting Intervals

You can specify the time after which an unacknowledged link state update should be transmitted on the interface. The integer value to specify retransmit intervals can range from 1 to 65,535 seconds.



Note

This value must be the same on the switches on both ends of the interface.

Send comments to nx5000-docfeedback@cisco.com

Configuring Retransmitting Intervals

To configure the FSPF retransmit time interval using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**.
You see the interface configuration in the Information pane.
 - Step 2** Click the **FSPF** tab.
You see the FSPF interface configuration in the Information pane as shown in [Figure 18-4](#).
 - Step 3** Double-click the ReTx Interval field and enter a value.
 - Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-

About Disabling FSPF for Specific Interfaces

You can disable the FSPF protocol for selected interfaces. By default, FSPF is enabled on all E ports and TE ports. This default can be disabled by setting the interface as passive.



Note

FSPF must be enabled at both ends of the interface for the protocol to work.

Disabling FSPF for Specific Interfaces

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.

To disable FSPF for a specific interface using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**.
You see the interface configuration in the Information pane.
 - Step 2** Click the **FSPF** tab.
You see the FSPF interface configuration in the Information pane as shown in [Figure 18-4](#).
 - Step 3** In the Admin Status drop-down list, choose **down**.
 - Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-

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.

Send comments to nx5000-docfeedback@cisco.com

Displaying the FSPF Database

The FSPF database for a specified VSAN includes the following information:

- Link State Record (LSR) type
- Domain ID of the LSR owner
- Domain ID of the advertising router
- LSR age
- LSR incarnation member
- Number of links

To display the FSPF database using Device Manager, perform this task:

Step 1 Choose **FC > Advanced > FSPF**.

You see the FSPF dialog box as shown in [Figure 18-5](#).

Figure 18-5 FSPF Dialog Box in Device Manager

VSAN Id	Admin Status	Oper Status	Set To Default?	Region Id	Domain Id	SPF HoldTime	SPF Delay	LSR Min Arrival (ms)	LSR Min Interval (ms)	LSR Refresh Time (min)	LSR Max Age (min)	CreateTime	CheckSum
1	up	up	<input type="checkbox"/>		0.0e67(103)	0	0	1000	2000	30	60	2007/04/09-19:14:47	351854
2	up	up	<input type="checkbox"/>		0.0eef(239)	0	0	1000	2000	30	60	2007/04/09-19:14:47	328940
3	up	up	<input type="checkbox"/>		0.0e2(2)	0	0	1000	2000	30	60	2007/04/09-19:14:47	192890
444	up	up	<input type="checkbox"/>		0.0e11(17)	0	0	1000	2000	30	60	2007/04/09-19:14:47	413687
551	up	up	<input type="checkbox"/>		0.0e3(227)	0	0	1000	2000	30	60	2007/04/09-19:14:47	266935
666	up	up	<input type="checkbox"/>		0.0e16(27)	0	0	1000	2000	30	60	2007/04/09-19:14:47	363053
999	up	up	<input type="checkbox"/>		0.0e67(231)	0	0	1000	2000	30	60	2007/04/09-19:14:47	421291
9001	up	up	<input type="checkbox"/>		0.0eef(239)	0	0	1000	2000	30	60	2007/04/09-19:14:47	229951
9002	up	up	<input type="checkbox"/>		0.0eef(239)	0	0	1000	2000	30	60	2007/04/09-19:14:47	297989
9003	up	up	<input type="checkbox"/>		0.0eef(239)	0	0	1000	2000	30	60	2007/04/09-19:14:47	310734

Step 2 Click the **LSDB LSRs** tab.

You see the FSPF database information as shown in [Figure 18-6](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 18-6 FSPF Database Information in the LSDB LSRs Tab

VSAN Id, DomainId	AdvDomainId	Age	IncarnationNumber	CheckSum	Links	External
1, 0x42 (66)	0x67(103)	230	0x80000177	0x1d5f	5	true
1, 0x61 (97)	0x61(97)	1253	0x800000d3	0xd50d	4	false
1, 0x62 (98)	0x62(98)	1262	0x800000d8	0x2a97	4	false
1, 0x63 (99)	0x63(99)	237	0x800000d8	0xcf4	9	false
1, 0x64 (100)	0x64(100)	836	0x800000d9	0xa8ed	10	false
1, 0x65 (101)	0x65(101)	831	0x800000da	0x17ac	9	false
1, 0x66 (102)	0x66(102)	831	0x800000d0	0xa391	3	false
1, 0x67 (103)	0x67(103)	830	0x800000e6	0x36d	15	false
1, 0x68 (104)	0x68(104)	1181	0x800000dd	0x9ee4	6	false
1, 0xd5 (213)	0xd5(213)	1013	0x80000901	0xe6f3	2	false
1, 0xd6 (214)	0xd6(214)	1447	0x8000090c	0xf821	3	false
2, 0x1 (1)	0x1(1)	1257	0x80000936	0x45bb	4	false
2, 0x4 (4)	0x4(4)	1191	0x80000a1c	0x615a	2	false

Step 3 Click **Close** to close the dialog box.

Viewing FSPF Statistics

To view FSPF statistics using Fabric Manager, perform this task:

- Step 1 Expand a Fabric, expand a VSAN, and then choose **FSPF** in the Logical Domains pane. You see the FSPF configuration dialog box.
- Step 2 Click the **Statistics** tab. You see the FSPF VSAN statistics in the Information pane as shown in Figure 18-7.

Figure 18-7 FSPF VSAN Statistics

Switch	Spt	Computations	Error Rcv	Errors	Checksum	LSU Rcv	LSU Tx	LSU Rcv Tx	LSA Rcv	LSA Tx	Hello Rcv	Hello Tx	Count
sw172-22-46-220		143	17	0	616	2136	8	2126	606	37223	37240	12	

Step 3 Click the **Interface Statistics** tab.

Send comments to nx5000-docfeedback@cisco.com

You see the FSPF interface statistics in the Information pane.

FSPF Routes

FSPF routes traffic across the fabric, based on entries in the FSPF database. These routes can be learned dynamically, or configured statically.


This section includes the following topics:

- [About Fibre Channel Routes, page 18-11](#)
- [Configuring Fibre Channel Routes, page 18-11](#)

About Fibre Channel Routes

Each port implements forwarding logic, which forwards frames based on its FC ID. Using the FC ID for the specified interface and domain, you can configure the specified route (for example, FC ID 111211 and domain ID 3) in the switch with domain ID 1 (see [Figure 18-8](#)).

Figure 18-8 Fibre Channel Routes



Configuring Fibre Channel Routes


If you disable FSPF, you can manually configure a Fibre Channel route. To configure a Fibre Channel route using Device Manager, perform this task:

Step 1 Click **FC > Advanced > Routes**.

You see the FC Static Route Configuration dialog box as shown in [Figure 18-9](#).

Send comments to nx5000-docfeedback@cisco.com


Figure 18-9 Fibre Channel Static Route Configuration Dialog Box



Step 2 Click **Create** to create a static route.

You see the Create Route dialog box as shown in [Figure 18-10](#).

Figure 18-10 Create Fibre Channel Route Dialog Box



Step 3 Choose the VSAN ID that for which you are configuring this route.

Step 4 Fill in the destination address and destination mask for the device you are configuring a route.

Step 5 Choose the interface that you want to use to reach this destination.

Step 6 Choose the next hop domain ID and route metric.

Step 7 Check either the **local** or **remote** radio button.

Step 8 Click **Create** to save these changes, or click **Close** to discard any unsaved changes.

In-Order Delivery

In-order delivery (IOD) of data frames guarantees frame delivery to a destination in the same order that they were sent by the originator.

Some Fibre Channel protocols or applications cannot handle out-of-order frame delivery. In these cases, switches in the Cisco Nexus 5000 Series preserve frame ordering in the frame flow. The source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) identify the flow of the frame.

On a switch with IOD enabled, all frames received by a specific ingress port and destined to a certain egress port are always delivered in the same order in which they were received.

Send comments to nx5000-docfeedback@cisco.com

Use IOD only if your environment cannot support out-of-order frame delivery.



Tip

If you enable the in-order delivery feature, the graceful shutdown feature is not implemented.


This section includes the following topics:

- [About Reordering Network Frames, page 18-13](#)
- [About Reordering SAN Port Channel Frames, page 18-13](#)
- [About Enabling In-Order Delivery, page 18-14](#)
- [Enabling In-Order Delivery Globally, page 18-15](#)
- [Enabling In-Order Delivery for a VSAN, page 18-15](#)
- [Configuring the Drop Latency Time, page 18-15](#)

About Reordering Network Frames

When you experience a route change in the network, the new selected path may be faster or less congested than the old route.

Figure 18-11 Route Change Delivery



In [Figure 18-11](#), the new path from Switch 1 to Switch 4 is faster. In this scenario, Frame 3 and Frame 4 may be delivered before Frame 1 and Frame 2.

If the in-order guarantee feature is enabled, the frames within the network are delivered as follows:


- Frames in the network are delivered in the order in which they are transmitted.
- Frames that cannot be delivered in order within the network latency drop period are dropped inside the network.

About Reordering SAN Port Channel Frames

When a link change occurs in a SAN port channel, the frames for the same exchange or the same flow can switch from one path to another faster path.

Send comments to nx5000-docfeedback@cisco.com

Figure 18-12 Link Congestion Delivery



In [Figure 18-12](#), the port of the old path (red dot) is congested. In this scenario, Frame 3 and Frame 4 can be delivered before Frame 1 and Frame 2.

When the in-order delivery feature is enabled and a port channel link change occurs, the frames crossing the SAN port channel are delivered as follows:

- Frames using the old path are delivered before new frames are accepted.
- The new frames are delivered through the new path after the network latency drop period has elapsed and all old frames are flushed.

Frames that cannot be delivered in order through the old path within the network latency drop period are dropped. See the [“Configuring the Drop Latency Time”](#) section on page 18-15.

About Enabling In-Order Delivery

You can enable the in-order delivery feature for a specific VSAN or for the entire switch. By default, in-order delivery is disabled on switches in the Cisco Nexus 5000 Series.



Tip

We recommend that you only enable this feature when devices that cannot handle any out-of-order frames are present in the switch. Load-balancing algorithms within the Cisco Nexus 5000 Series switch 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 the in-order delivery feature is enabled, the recovery will be delayed because of an intentional pausing of fabric forwarding to purge the fabric of resident frames that could potentially be forwarded out-of-order.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Enabling In-Order Delivery Globally

To ensure that the in-order delivery parameters are uniform across all VSANs on the switch, enable in-order delivery globally.

Only enable in-order delivery globally if this is a requirement across your entire fabric. Otherwise, enable IOD only for the VSANs that require this feature.

Enabling In-Order Delivery for a VSAN

When you create a VSAN, that VSAN automatically inherits the global in-order guarantee value. You can override this global value by enabling or disabling in-order guarantee for the new VSAN.

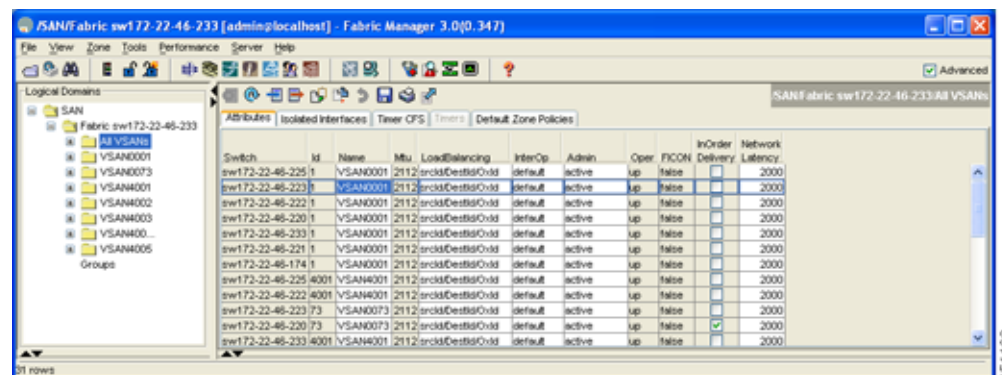
To use the lowest domain switch for the multicast tree computation using Fabric Manager, perform this task:

Step 1 Expand a fabric and then choose **All VSANS**.

Step 2 Click the **Attributes** tab.

You see the general VSAN attributes in the Information pane as shown in [Figure 18-13](#).

Figure 18-13 General VSAN Attributes



Step 3 Check the **InOrder Delivery** check box to enable IOD for the switch.

Step 4 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

Configuring the Drop Latency Time

You can change the default latency time for either the entire switch or a specified VSAN in a switch.

To configure the drop latency time for a switch using Fabric Manager, perform this task:

Step 1 Expand a fabric and then choose **All VSANS**.

You see the VSAN configuration in the Information pane.

Step 2 Click the **Attributes** tab.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

You see the general VSAN attributes in the Information pane as shown in [Figure 18-14](#).

Figure 18-14 General VSAN Attributes



Step 3 Double-click the Network Latency field and change the value.

Step 4 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

Default Settings

[Table 18-2](#) lists the default settings for FSPF features.

Table 18-2 Default FSPF Settings

Parameters	Default
FSPF	Enabled on all E ports and TE ports.
SPF computation	Dynamic.
SPF hold time	0.
Backbone region	0.
Acknowledgment interval (RxmtInterval)	5 seconds.
Refresh time (LSRefreshTime)	30 minutes.
Maximum age (MaxAge)	60 minutes.
Hello interval	20 seconds.
Dead interval	80 seconds.
Distribution tree information	Derived from the principal switch (root node).
Routing table	FSPF stores up to 16 equal cost paths to a given destination.
Load balancing	Based on destination ID and source ID on different, equal cost paths.
In-order delivery	Disabled.
Drop latency	Disabled.

Send comments to nx5000-docfeedback@cisco.com

Table 18-2 *Default FSPF Settings (continued)*

Parameters	Default
Static route cost	If the cost (metric) of the route is not specified, the default is 10.
Remote destination switch	If the remote destination switch is not specified, the default is direct.
Multicast routing	Uses the principal switch to compute the multicast tree.

Send comments to nx5000-docfeedback@cisco.com



CHAPTER 19

Managing FLOGI, Name Server, FDMI, and RSCN Databases

This chapter describes the fabric login (FLOGI) database, the name server features, the Fabric-Device Management Interface (FDMI), and Registered State Change Notification (RSCN) information provided in Cisco Nexus 5000 Series switches. It includes the following sections:

- [Information About Fabric Login, page 19-1](#)
- [Name Server Proxy, page 19-2](#)
- [FDMI, page 19-4](#)
- [Displaying FDMI, page 19-4](#)
- [RSCN, page 19-5](#)
- [Default Settings, page 19-8](#)

Information About Fabric Login

In a Fibre Channel fabric, each host or disk requires an FC ID. If the required device is displayed in the FLOGI table, the fabric login is successful. Examine the FLOGI database on a switch that is directly connected to the host HBA and connected ports. See the [“Default Company ID List” section on page 22-7](#) and the [“Switch Interoperability” section on page 22-7](#).

To verify that a storage device is in the fabric login (FLOGI) table using Fabric Manager, perform this task:

-
- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**.
You see the interface configuration in the Information pane.
- Step 2** Click the **FLOGI** tab.
You see all end devices that are logged into the fabric as shown in [Figure 19-1](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 19-1 FLOGI Physical Interfaces

Switch	Interface, VSAN Id	FcId	PortName	NodeName	Version	CoS	Class 2 RxDataSize	Class 2 SeqDeliv	Class 3 RxDataSize	Class 3 SeqDeliv
sw172-22-46-224	Fc1/6, 4001	isa0197	Seagate 22:00:00:20:37:73:de:d5	Seagate 20:00:00:20:37:73:de:d5	32.3		0	false	2112	true
sw172-22-46-224	Fc1/6, 4001	isa0198	Seagate 22:00:00:20:37:46:56:52	Seagate 20:00:00:20:37:46:56:52	32.3		0	false	2112	true
sw172-22-46-224	Fc1/6, 4001	isa019f	Seagate 22:00:00:20:37:46:39:14	Seagate 20:00:00:20:37:46:39:14	32.3		0	false	2112	true
sw172-22-46-224	Fc1/6, 4001	isa01a2	Seagate 22:00:00:20:37:5b:b1:8e	Seagate 20:00:00:20:37:5b:b1:8e	32.3		0	false	2112	true
sw172-22-46-224	Fc1/6, 4001	isa01a7	Seagate 22:00:00:20:37:5b:81:1b	Seagate 20:00:00:20:37:5b:81:1b	32.3		0	false	2112	true

Name Server Proxy

The name server functionality maintains a database containing the attributes for all hosts and storage devices in each VSAN. Name servers allow a database entry to be modified by a device that originally registered the information.

The proxy feature is useful when you need to modify (update or delete) the contents of a database entry that was previously registered by a different device.

This section includes the following topics:

- [About Registering Name Server Proxies, page 19-2](#)
- [Registering Name Server Proxies, page 19-2](#)
- [About Rejecting Duplicate pWWNs, page 19-3](#)
- [Rejecting Duplicate pWWNs, page 19-3](#)
- [About Name Server Database Entries, page 19-3](#)
- [Viewing Name Server Database Entries, page 19-3](#)

About Registering Name Server Proxies

All name server registration requests come from the same port whose parameter is registered or changed. If it does not, then the request is rejected.

This authorization enables WWNs to register specific parameters for another node.


Registering Name Server Proxies

To register the name server proxy using Fabric Manager, perform this task:

- Step 1** Expand a fabric, expand a VSAN, and then choose **Advanced**.
You see the VSAN advanced configuration in the Information pane.
- Step 2** Click the **NS Proxies** tab.
You see the existing name server proxy for the selected VSAN as shown in [Figure 19-2](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 19-2 Name Server Proxies



- Step 3** Double-click the PortName field and enter a new value to register a new name server proxy.
- Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to cancel any unsaved changes.

About Rejecting Duplicate pWWNs

You can prevent malicious or accidental log in using another device's pWWN. These pWWNs are allowed to log in to the fabric and replace the first device in the name server database.

Rejecting Duplicate pWWNs

To reject duplicate pWWNs, see the *Cisco Cisco Nexus 5000 Series CLI Configuration Guide*.

About Name Server Database Entries

The name server stores name entries for all hosts in the FCNS database. The name server permits an Nx port to register attributes during a PLOGI (to the name server) to obtain attributes of other hosts. These attributes are deregistered when the Nx port logs out either explicitly or implicitly.

In a multiswitch fabric configuration, the name server instances running on each switch shares information in a distributed database. One instance of the name server process runs on each switch.

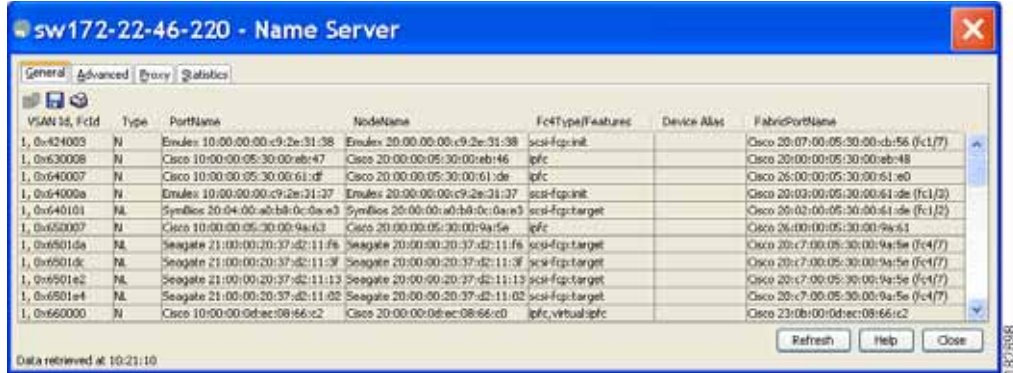
Viewing Name Server Database Entries

To view the name server database using Device Manager, perform this task:

- Step 1** Click **FC > Name Server**.
- You see the Name Server dialog box as shown in [Figure 19-3](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 19-3 Name Server Dialog Box



The General tab is the default tab; you see the name server database.

- Step 2** Click the **Statistics** tab.
- You see the name server statistics.
- Step 3** Click **Close** to close the dialog box.

FDMI

Cisco Nexus 5000 Series switches provide support for the Fabric-Device Management Interface (FDMI) functionality, as described in the FC-GS-4 standard. FDMI enables management of devices such as Fibre Channel host bus adapters (HBAs) through in-band communications. This addition complements the existing Fibre Channel name server and management server functions.

Using the FDMI functionality, the switch software can extract the following management information about attached HBAs and host operating systems without installing proprietary host agents:

- Manufacturer, model, and serial number
- Node name and node symbolic name
- Hardware, driver, and firmware versions
- Host operating system (OS) name and version number

All FDMI entries are stored in persistent storage and are retrieved when the FDMI process is started.

Displaying FDMI

To display the FDMI database information using Device Manager, perform this task:

- Step 1** Click **FC > Advanced > FDMI**
- You see the FDMI dialog box.
- Step 2** Click the **HBA** tab, the **Versions** tab or the **Targets** tab.

Send comments to nx5000-docfeedback@cisco.com

RSCN

The Registered State Change Notification (RSCN) is a Fibre Channel service that informs hosts about changes in the fabric. Hosts can receive this information by registering with the fabric controller (through SCR). These notifications provide a timely indication of one or more of the following events:

- Disks joining or leaving the fabric
- A name server registration change
- A new zone enforcement
- IP address change
- Any other similar event that affects the operation of the host

This section includes the following topics:

- [About RSCN Information, page 19-5](#)
- [Displaying RSCN Information, page 19-5](#)
- [About the multi-pid Option, page 19-6](#)
- [Configuring the multi-pid Option, page 19-6](#)
- [Configuring the multi-pid Option, page 19-6](#)
- [Configuring the multi-pid Option, page 19-6](#)
- [Clearing RSCN Statistics, page 19-7](#)
- [RSCN Timer Configuration Distribution Using CFS, page 19-7](#)
- [Configuring the RSCN Timer with CFS, page 19-8](#)

About RSCN Information

A switch RSCN (SW-RSCN) is sent to registered hosts and to all reachable switches in the fabric.



Note

The switch sends an RSCN to notify registered nodes that a change has occurred. It is up to the nodes to query the name server again to obtain the new information. The details of the changed information are not delivered by the switch in the RSCN sent to the nodes.

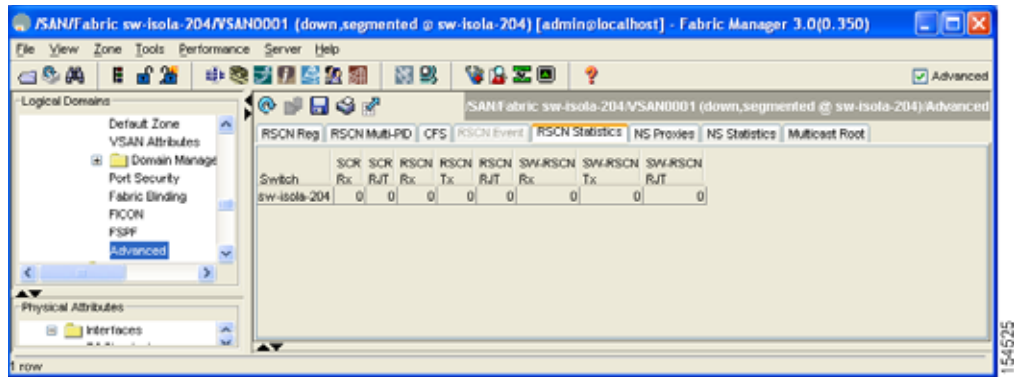
Displaying RSCN Information

To display RSCN information using Fabric Manager, perform this task:

- Step 1** Expand a fabric, expand a VSAN, and then choose **Advanced**.
You see the VSAN advanced configuration in the Information pane.
- Step 2** Click the **RSCN Reg** tab or the **RSCN Statistics** tab (see [Figure 19-4](#)).
You see the RSCN Statistics pane as shown in [Figure 19-4](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 19-4 RSCN Statistics



About the multi-pid Option

If the RSCN **multi-pid** option is enabled, then RSCNs generated to the registered Nx ports may contain more than one affected port IDs. In this case, zoning rules are applied before putting the multiple affected port IDs together in a single RSCN. By enabling this option, you can reduce the number of RSCNs. For example, you have two disks (D1, D2) and a host (H) connected to switch 1. Host H is registered to receive RSCNs. D1, D2 and H belong to the same zone. If disks D1 and D2 are online at the same time, one of the following actions applies:

- The **multi-pid** option is disabled on switch 1— Two RSCNs are generated to host H: one for the disk D1 and another for disk D2.
- The **multi-pid** option is enabled on switch 1—A single RSCN is generated to host H, and the RSCN payload lists the affected port IDs (in this case, both D1 and D2).



Note

Some Nx ports may not support multi-pid RSCN payloads. If so, disable the RSCN **multi-pid** option.


Configuring the multi-pid Option

To configure the **multi-pid** option using Fabric Manager, perform this task:

- Step 1** Expand a fabric, expand a VSAN, and then choose **Advanced**.
You see the VSAN advanced configuration in the Information pane.
- Step 2** Click the **RSCN Multi-PID** tab.
You see the RSCN Multi-PID pane as shown in [Figure 19-5](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 19-5 RSCN Multi-PID



Step 3 Check the **Enable** check box.

Step 4 Click **Apply Changes** to save these changes, or click **Undo Changes** to cancel any unsaved changes.

Clearing RSCN Statistics

You can clear the counters and later view the counters for a different set of events. For example, you can keep track of how many RSCNs or SW-RSCNs are generated on a particular event (such as ONLINE or OFFLINE events). You can use these statistics to monitor responses for each event in the VSAN.

To clear the RSCN statistics for the specified VSAN, see the *Cisco Cisco Nexus 5000 Series CLI Configuration Guide*.

RSCN Timer Configuration Distribution Using CFS

Because the timeout value for each switch is configured manually, a misconfiguration occurs when different switches time out at different times. This means different N-ports in a network can receive RSCNs at different times. Cisco Fabric Services (CFS) alleviates this situation by automatically distributing configuration information to all switches in a fabric. This also reduces the number of SW-RSCNs.

RSCN supports two modes, distributed and nondistributed. In distributed mode, RSCN uses CFS to distribute configuration to all switches in the fabric. In nondistributed mode, only the configuration commands on the local switch are affected.



Note

All configuration commands are not distributed. Only the **rscn event-tov tov vsan vsan** command is distributed.

The RSCN timer is registered with CFS during initialization and switchover. For high availability, if the RSCN timer distribution crashes and restarts or a switchover occurs, it resumes normal functionality from the state prior to the crash or switchover.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

**Note**

Before performing a downgrade, make sure that you revert the RSCN timer value in your network to the default value. Failure to do so will disable the links across your VSANs and other devices.

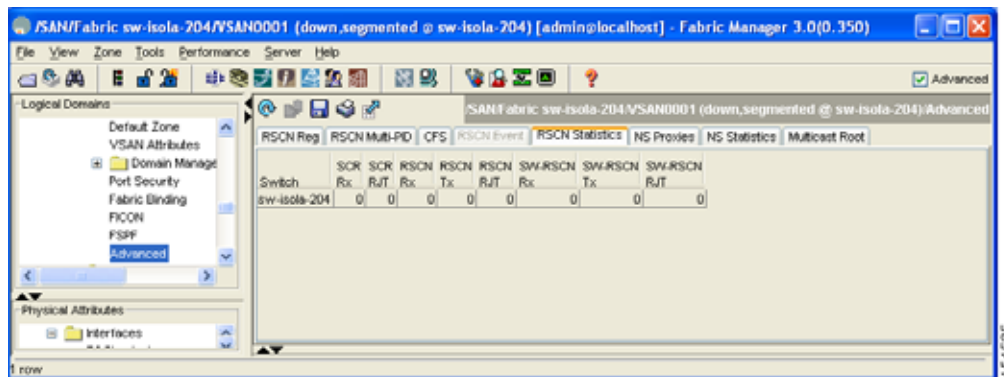
Compatibility across various software releases during an upgrade or downgrade is supported by **conf-check** provided by CFS. You are required to disable RSCN timer distribution support before you downgrade. By default, the RSCN timer distribution capability is disabled.

Configuring the RSCN Timer with CFS

To configure the RSCN timer with CFS using Fabric Manager, perform this task:

- Step 1** Expand a fabric, expand a VSAN, and then choose **Advanced** in the Logical Domains pane.
- Step 2** Click the **RSCN Event** tab.
- You see the VSAN advanced configuration in the Information pane as shown in [Figure 19-6](#).

Figure 19-6 VSAN Advanced Configuration



- Step 3** Double-click the **TimeOut** value to change the value (in milliseconds) for the selected VSAN.
- Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to cancel any unsaved changes.

In this example the event time-out value is set to 300 milliseconds for VSAN 12.

Default Settings

[Table 19-1](#) lists the default settings for RSCN.

Table 19-1 Default RSCN Settings

Parameters	Default
RSCN timer value	2000 milliseconds for Fibre Channel VSANs
RSCN timer configuration distribution	Disabled

Send comments to nx5000-docfeedback@cisco.com

Send comments to nx5000-docfeedback@cisco.com



CHAPTER 20

Configuring SPAN

The Switched Port Analyzer (SPAN) feature (sometimes called port mirroring or port monitoring) selects network traffic for analysis by a network analyzer. The network analyzer can be a Cisco SwitchProbe, a Fibre Channel Analyzer, or other Remote Monitoring (RMON) probes.

This section includes the following topics:

- [Information About SPAN Sources, page 20-1](#)
- [Information About SPAN Destinations, page 20-2](#)
- [Configuring SPAN, page 20-3](#)
- [Default SPAN Settings, page 20-5](#)

Information About SPAN Sources

SPAN sources refer to the interfaces from which traffic can be monitored. The Cisco Nexus 5000 Series switch supports Ethernet, virtual Ethernet, Fibre Channel, virtual Fibre Channel, Port Channels, SAN port channels, VLANs, and VSANs as SPAN sources. In the case of VLANs or VSANs, all supported interfaces in the specified VLAN or VSAN are included as SPAN sources. You can choose the SPAN traffic in the ingress direction, the egress direction, or both directions for Ethernet, virtual Ethernet, Fibre Channel, and virtual Fibre Channel source interfaces:

- Ingress source (Rx)—Traffic entering the switch through this source interface is copied to the SPAN destination port.
- Egress source (Tx)—Traffic exiting the switch through this source interface is copied to the SPAN destination port.



Note

Device Manager does not support the configuration of Ethernet or virtual Ethernet interfaces as source ports.

Characteristics of Source Ports

A source port, also called a monitored port, is a switched interface that you monitor for network traffic analysis. The switch supports any number of ingress source ports (up to the maximum number of available ports on the switch) and any number of source VLANs or VSANs.

A source port has these characteristics:

Send comments to nx5000-docfeedback@cisco.com

- Can be any port type: Ethernet, virtual Ethernet, Fibre Channel, virtual Fibre Channel, Port Channel, SAN-Port Channel, VLAN, and VSAN.
- Cannot be monitored in multiple SPAN sessions.
- Cannot be a destination port.
- Each source port can be configured with a direction (ingress, egress, or both) to monitor. For VLAN, VSAN, Port Channel, and SAN-Port Channel sources, the monitored direction can only be ingress and applies to all physical ports in the group.
- Source ports can be in the same or different VLANs or VSANs.
- For VLAN or VSAN SPAN sources, all active ports in the source VLAN or VSAN are included as source ports.
- The switch supports a maximum of two egress SPAN source ports.

Information About SPAN Destinations

SPAN destinations refer to the interfaces that monitors source interfaces. The Cisco Nexus 5000 Series switch supports Ethernet and Fibre Channel interfaces as SPAN destinations.



Note

Device Manager does not support the configuration of Ethernet interfaces as destination ports.

Source SPAN	Dest SPAN
Ethernet	Ethernet
Fibre Channel	Fibre Channel
Fibre Channel	Ethernet (FCoE)
Virtual Ethernet	Ethernet
Virtual Fibre Channel	Fibre Channel
Virtual Fibre Channel	Ethernet (FCoE)

Characteristics of Destination Ports

Each local SPAN session must have a destination port (also called a monitoring port) that receives a copy of traffic from the source ports, VLANs, or VSANs. A destination port has these characteristics:

- Can be any physical port: Ethernet, Ethernet (FCoE), or Fibre Channel. Virtual Ethernet and virtual Fibre Channel ports cannot be destination ports.
- Cannot be a source port.
- Cannot be a Port Channel or SAN-Port Channel group.
- Does not participate in spanning tree while the SPAN session is active.
- Is excluded from the source list and is not monitored, if it belongs to a source VLAN of any SPAN session.

Send comments to nx5000-docfeedback@cisco.com

- Receives copies of sent and received traffic for all monitored source ports. If a destination port is oversubscribed, it can become congested. This congestion can affect traffic forwarding on one or more of the source ports.

Configuring SPAN

You can configure a SPAN session to duplicate packets from source ports to the specified destination ports on the switch. This section includes the following topics:

- [Configuring SPAN Using Device Manager, page 20-3](#)
- [Creating SPAN Sessions Using Device Manager, page 20-3](#)
- [Editing SPAN Sources Using Device Manager, page 20-4](#)
- [Deleting SPAN Sessions Using Device Manager, page 20-5](#)

Configuring SPAN Using Device Manager

To monitor network traffic using SD ports, perform this task:

-
- | | |
|--------|---|
| Step 1 | Configure the SD port. |
| Step 2 | Attach the SD port to a specific SPAN session. |
| Step 3 | Monitor network traffic by adding source interfaces to the session. |
-

To configure an SD port for SPAN monitoring using Device Manager, perform this task:

-
- | | |
|--------|--|
| Step 1 | Right-click the port that you want to configure and click Configure .
You see the general port configuration dialog box. |
| Step 2 | Under Mode, choose SD . |
| Step 3 | Click Apply to accept the change. |
| Step 4 | Close the dialog box. |
-


Creating SPAN Sessions Using Device Manager

To create SPAN sessions using Device Manager, perform this task:

-
- | | |
|--------|--|
| Step 1 | Choose Interface > SPAN .
You see the SPAN dialog box. |
| Step 2 | Click the Sessions tab. |
| Step 3 | Click Create .
You see the Create SPAN Sessions dialog box as shown in Figure 20-1 . |

Send comments to nx5000-docfeedback@cisco.com

Figure 20-1 Create SPAN Sessions Dialog Box



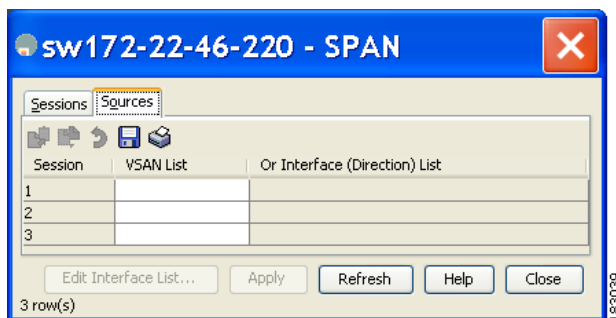
- Step 4 Choose the session ID (from 1-16) using the up or down arrows and click **Create**.
- Step 5 Repeat Step 4 for each session that you want to create.
- Step 6 Enter the destination interface in the Dest Interface field for the appropriate session.
- Step 7 Enter the filter VSAN list in the Filter VSAN List field for the appropriate session.
- Step 8 Choose **active** or in **active** admin status in the Admin drop-down list.
- Step 9 Click **Apply** to save your changes.
- Step 10 Close the two dialog boxes.

Editing SPAN Sources Using Device Manager

To edit a SPAN source using Device Manager, perform this task:

- Step 1 Choose **Interface > SPAN**.
You see the SPAN dialog box.
- Step 2 Click the **Sources** tab.
You see the SPAN Sources dialog box as shown in [Figure 20-2](#).

Figure 20-2 SPAN Sources Tab




- Step 3 Enter the VSAN list name in the VSAN List field.
- Step 4 Click **Edit Interface List**.
You see the Source Interfaces dialog box.
- Step 5 Click **Create**.

Send comments to nx5000-docfeedback@cisco.com

You see the Source Interfaces Interface Sources dialog box as shown in [Figure 20-3](#).

Figure 20-3 Source Interfaces Interface Sources Dialog Box



- Step 6** Click the browse button to display the list of available FC ports.
- Step 7** Choose a port and click **OK**.
- Step 8** Check the direction (**receive** or **transmit**) that you want.
- Step 9** Click **Create** to create the FC interface source.
- Step 10** Click **Close** in each of the three open dialog boxes.

Deleting SPAN Sessions Using Device Manager

To delete a SPAN session using Device Manager, perform this task:

- Step 1** Choose **Interface > SPAN**.
You see the SPAN dialog box.
- Step 2** Click the **Sessions** tab.
- Step 3** Click the SPAN session that you want to delete.
- Step 4** Click **Delete**.
The SPAN session is deleted.
- Step 5** Close the dialog box.

Default SPAN Settings

[Table 20-1](#) lists the default settings for SPAN parameters.

Table 20-1 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.

Send comments to nx5000-docfeedback@cisco.com

Table 20-1 Default SPAN Configuration Parameters (continued)

Parameters	Default
Encapsulation	Disabled.
SD port	Output frame format is Fibre Channel.



Discovering SCSI Targets

This chapter describes the SCSI LUN discovery feature provided in switches in the Cisco Nexus 5000 Series. It includes the following sections:

- [Information About SCSI LUN Discovery, page 21-1](#)
- [Displaying SCSI LUN Information, page 21-3](#)

Information 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 that a Network Management System (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 Nexus 5000 Series.

This section includes the following topics:

- [About Starting SCSI LUN Discovery, page 21-1](#)
- [Starting SCSI LUN Discovery, page 21-2](#)
- [About Initiating Customized Discovery, page 21-2](#)
- [Initiating Customized Discovery, page 21-2](#)

About Starting SCSI LUN Discovery

SCSI LUN discovery is done on demand.

Only Nx ports that are present in the name server database and that are registered as FC4 Type = SCSI_FCP are discovered.

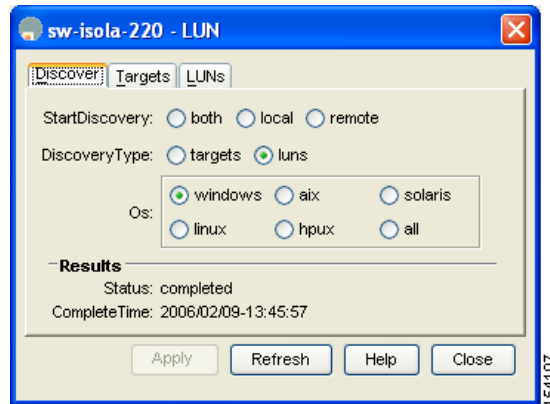
Send comments to nx5000-docfeedback@cisco.com

Starting SCSI LUN Discovery

To begin SCSI LUN discovery using Device Manager, perform this task:

-
- Step 1** Choose **FC > Advanced > LUNs**.
You see the LUN Configuration dialog box as shown in [Figure 21-1](#).

Figure 21-1 LUN Configuration Dialog Box



- Step 2** Set StartDiscovery to **local**, **remote** or **both**.
Step 3 Choose the **DiscoveryType** and **OS**.
Step 4 Click **Apply** to begin discovery.
-

About Initiating Customized Discovery

Customized discovery consists of a list of VSAN and domain pairs that are selectively configured to initiate a discovery. The domain ID is a number from 0 to 255 in decimal or a number from 0x0 to 0xFF in hex.

Initiating Customized Discovery

To initiate a customized discovery using Device Manager, perform this task:

-
- Step 1** Click the VSAN drop-down menu and choose the VSAN in which you want to initiate a customized discovery.
Step 2 Choose **FC > Advanced > LUNs**.
You see the LUN Configuration dialog box.
Step 3 Set StartDiscovery to **local**, **remote** or **both**.
Step 4 Fill in the DiscoveryType and OS fields.

Send comments to nx5000-docfeedback@cisco.com

Step 5 Click **Apply** to begin discovery.

Displaying SCSI LUN Information

To display the results of the discovery using Device Manager, perform this task:

-
- Step 1** Choose **FC > Advanced > LUNs**
You see the LUN Configuration dialog box.
- Step 2** Click the **LUN** tab or the **Targets** tab.
-

Send comments to nx5000-docfeedback@cisco.com



Advanced Features and Concepts

This chapter describes the advanced Fibre Channel features provided in Cisco Nexus 5000 Series switches. It includes the following sections:

- [Fibre Channel Timeout Values, page 22-1](#)
- [World Wide Names, page 22-5](#)
- [FC ID Allocation for HBAs, page 22-6](#)
- [Switch Interoperability, page 22-7](#)
- [Default Settings, page 22-12](#)

Fibre Channel Timeout Values

You can modify Fibre Channel protocol-related timer values for the switch by configuring the following timeout values (TOVs):

- Distributed services TOV (D_S_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 5,000 milliseconds.
- Error detect TOV (E_D_TOV)—The valid range is from 1,000 to 10,000 milliseconds. The default is 2,000 milliseconds. This value is matched with the other end during port initialization.
- Resource allocation TOV (R_A_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 10,000 milliseconds. This value is matched with the other end during port initialization.



Note

The fabric stability TOV (F_S_TOV) constant cannot be configured.

This section includes the following topics:

- [Timer Configuration Across All VSANs, page 22-2](#)
- [Timer Configuration Per-VSAN, page 22-3](#)
- [About fctimer Distribution, page 22-4](#)
- [Enabling or Disabling fctimer Distribution, page 22-4](#)
- [Database Merge Guidelines, page 22-4](#)

Send comments to nx5000-docfeedback@cisco.com

Timer Configuration Across All VSANs

You can modify Fibre Channel protocol related timer values for the switch.



Caution

The D_S_TOV, E_D_TOV, and R_A_TOV values cannot be globally changed unless all VSANs in the switch are suspended.



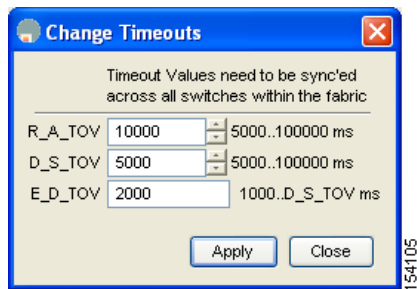
Note

If a VSAN is not specified when you change the timer value, the changed value is applied to all VSANs in the switch.

To configure timers in Fabric Manager, perform this task:

-
- Step 1** Expand **Switches > FC Services**, and then choose **Timers & Policies** in the Physical Attributes pane. You see the timers for multiple switches in the Information pane.
- Step 2** Click the **Change Timeouts** button to configure the timeout values. You see the dialog box as shown in [Figure 22-1](#).

Figure 22-1 *Configure Timers in Fabric Manager*




To configure timers in Device Manager, perform this task:

-
- Step 1** Choose **FC > Advanced > Timers/Policies**. You see the timers for a single switch in the dialog box as shown in [Figure 22-2](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 22-2 Configure Timers in Device Manager



Timer Configuration Per-VSAN

You can also issue the `ftimer` for a specified VSAN to configure different TOV values for VSANs with special links such as Fibre Channel. You can configure different `E_D_TOV`, `R_A_TOV`, and `D_S_TOV` values for individual VSANs. Active VSANs are suspended and activated when their timer values are changed.



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.

To configure per-VSAN Fiber Channel timers using Device Manager, perform this task:

Step 1 Choose **FC > Advanced > VSAN Timers**.

You see the VSANs Timer dialog box as shown in [Figure 22-3](#).

Figure 22-3 VSAN Timers in Device Manager

VSAN Id	R_A_TOV	D_S_TOV	E_D_TOV	NetworkDropLatency (ms)
1	10000	5000	2000	2000
2	10000	5000	2000	2000
3	10000	5000	2000	2000
444	10000	5000	2000	2000
501	10000	5000	2000	2000
666	10000	5000	2000	2000
999	10000	5000	2000	2000
4001	10000	5000	2000	2000
4002	10000	5000	2000	2000
4003	10000	5000	2000	2001

Step 2 Fill in the timer values that you want to configure.

Send comments to nx5000-docfeedback@cisco.com

Step 3 Click **Apply** to save these changes.

About fctimer Distribution

You can enable per-VSAN fctimer fabric distribution for all Cisco SAN switches in the fabric. When you perform fctimer configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you enter the first configuration command after you enabled distribution in a switch. The fctimer application uses the effective and pending database model to store or commit the commands based on your configuration.

See [Chapter 7, “Using Cisco Fabric Services,”](#) for more information on the CFS application.

Enabling or Disabling fctimer Distribution

To enable and distribute fctimer configuration changes using Device Manager, perform this task:

- Step 1** Choose **FC > Advanced > VSAN Timers**.
You see the VSANs Timer dialog box as shown in [Figure 22-3](#).
 - Step 2** Fill in the timer values that you want to configure.
 - Step 3** Click **Apply** to save these changes.
 - Step 4** Choose **commit** from the CFS drop-down list to distribute these changes or choose **abort** from the CFS drop-down list to discard any unsaved changes.
-

When you commit the fctimer configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. When you commit the fctimer configuration changes without implementing the session feature, the fctimer configurations are distributed to all the switches in the physical fabric.

Database Merge Guidelines

When merging two fabrics, follow these guidelines:

- Be aware of the following merge conditions:
 - The merge protocol is not implemented for distribution of the fctimer values. You must manually merge the fctimer values when a fabric is merged.
 - The per-VSAN fctimer configuration is distributed in the physical fabric.
 - The fctimer configuration is only applied to those switches containing the VSAN with a modified fctimer value.
 - The global fctimer values are not distributed.
- Do not configure global timer values when distribution is enabled.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)



Note

The number of pending fctimer configuration operations cannot be more than 15. After 15 operations, you must commit or abort the pending configurations before performing any more operations.

See the “[CFS Merge Support](#)” section on page 7-6 for additional information.

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.

Cisco Nexus 5000 Series switches support three network address authority (NAA) address formats (see [Table 22-1](#)).

Table 22-1 Standardized NAA WWN Formats

NAA Address	NAA Type	WWN Format	
IEEE 48-bit address	Type 1 = 0001b	000 0000 0000b	48-bit MAC address
IEEE extended	Type 2 = 0010b	Locally assigned	48-bit MAC address
IEEE registered	Type 5 = 0101b	IEEE company ID: 24 bits	VSID: 36 bits



Caution

Changes to the world-wide names should be made by an administrator or individual who is completely familiar with switch operations.

This section includes the following topics:

- [Verifying WWN Information, page 22-5](#)
- [Link Initialization WWN Usage, page 22-5](#)
- [Configuring a Secondary MAC Address, page 22-6](#)

Verifying WWN Information

To display WWN information using Device Manager, choose **FC > Advanced > WWN Manager**. You see the list of allocated WWNs.

Link Initialization WWN Usage

Exchange Link Protocol (ELP) and Exchange Fabric Protocol (EFP) use WWNs during link initialization. ELPs and EFPs both use the VSAN WWN by default during link initialization. However, the ELP usage changes based on the peer switch’s usage:

- If the peer switch ELP uses the switch WWN, then the local switch also uses the switch WWN.
- If the peer switch ELP uses the VSAN WWN, then the local switch also uses the VSAN WWN.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)


Configuring a Secondary MAC Address

To allocate secondary MAC addresses using Device Manager, perform this task:

Step 1 Choose **FC > Advanced > WWN Manager**.

You see the list of allocated WWNs as shown in [Figure 22-4](#).

Figure 22-4 Allocated World Wide Names in Device Manager



Step 2 Fill in the BaseMacAddress and MacAddressRange fields with the appropriate values.

Step 3 Click **Apply** to save these changes, or click **Close** to discard any unsaved changes.

FC ID Allocation for HBAs

Fibre Channel standards require a unique FC ID to be allocated to an N port attached to an F port in any switch. To conserve the number of FC IDs used, Cisco Nexus 5000 Series switches use a special allocation scheme.

Some HBAs do not discover targets that have FC IDs with the same domain and area. The switch software maintains a list of tested company IDs that do not exhibit this behavior. These HBAs are allocated with single FC IDs. If the HBA can discover targets within the same domain and area, a full area is allocated.

To allow further scalability for switches with numerous ports, the switch software maintains a list of HBAs that can discover targets within the same domain and area. Each HBA is identified by its company ID (also known as Organizational Unique Identifier, or OUI) used in the pWWN during a fabric log in. A full area is allocated to the N ports with company IDs that are listed and for the others, a single FC ID is allocated. Regardless of the type (whole area or single) of FC ID allocated, the FC ID entries remain persistent.

This section includes the following topics:

- [Default Company ID List, page 22-7](#)
- [Verifying the Company ID Configuration, page 22-7](#)

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Default Company ID List

All Cisco Nexus 5000 Series switches contain a default list of company IDs that require area allocation. Using the company ID reduces the number of configured persistent FC ID entries. You can configure or modify these entries using the CLI.



Caution

Persistent entries take precedence over company ID configuration. If the HBA fails to discover a target, verify that the HBA and the target are connected to the same switch and have the same area in their FC IDs, then perform the following procedure:

1. Shut down the port connected to the HBA.
2. Clear the persistent FC ID entry.
3. Get the company ID from the port WWN.
4. Add the company ID to the list that requires area allocation.
5. Bring up the port.

The list of company IDs have the following characteristics:

- A persistent FC ID configuration always takes precedence over the list of company IDs. Even if the company ID is configured to receive an area, the persistent FC ID configuration results in the allocation of a single FC ID.
- New company IDs added to subsequent releases are automatically added to existing company IDs.
- The list of company IDs is saved as part of the running and saved configuration.
- The list of company IDs is used only when the fcinterop FC ID allocation scheme is in auto mode. By default, the interop FC ID allocation is set to auto, unless changed.



Tip

We recommend that you set the fcinterop FC ID allocation scheme to auto and use the company ID list and persistent FC ID configuration to manipulate the FC ID device allocation.

See the *Cisco Cisco Nexus 5000 Series CLI Configuration Guide* to change the FC ID allocation.

The following example adds a new company ID to the default list.

```
switch(config)# fcid-allocation area company-id 0x003223
```

Verifying the Company ID Configuration

To view the configured company IDs using Device Manager, choose **FC > Advanced > FcId Area Allocation**. You can implicitly derive the default entries shipped with a specific release by combining the list of Company IDs displayed without any identification with the list of deleted entries.

Some WWN formats do not support company IDs. In these cases, you may need to configure the FC ID persistent entry.

Switch Interoperability

Interoperability enables the products of multiple vendors to interwork with each other. Fibre Channel standards guide vendors towards common external Fibre Channel interfaces.

Send comments to nx5000-docfeedback@cisco.com

Not all vendors follow the standards in the same way, which results in the need for interoperability modes. This section briefly explains the basic concepts of these modes.

Each vendor has a regular mode and an equivalent interoperability mode, which specifically turns off advanced or proprietary features and provides the product with a standards-compliant implementation.

This section includes the following topics:

- [About Interop Mode, page 22-8](#)
- [Configuring Interop Mode 1, page 22-9](#)
- [Verifying Interoperating Status, page 22-11](#)

About Interop Mode

Cisco NX-OS software supports the following four interop modes:

- Mode 1—Standards-based interop mode that requires all other vendors in the fabric to be in interop mode.
- Mode 2—Brocade native mode (Core PID 0).
- Mode 3—Brocade native mode (Core PID 1).
- Mode 4—McData native mode.

For information about configuring interop modes 2, 3, and 4, see the [Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide](#).

[Table 22-2](#) lists the changes in switch operation when you enable interoperability mode. These changes are specific to Cisco Nexus 5000 Series switches while in interop mode.

Table 22-2 Changes in Switch Operation 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 to 127, to accommodate McData's nominal restriction to this same range. Domain IDs can either be static or preferred, which operate as follows: <ul style="list-style-type: none"> • Static: Cisco switches accept only one domain ID; if a switch does not get that domain ID it isolates itself from the fabric. • Preferred: If the switch does not get its requested domain ID, it accepts any assigned domain ID.
Timers	All Fibre Channel timers must be the same on all switches as these values are exchanged by E ports when establishing an ISL. The timers are F_S_TOV, D_S_TOV, E_D_TOV, and R_A_TOV.
F_S_TOV	Verify that the Fabric Stability Time Out Value timers match exactly.
D_S_TOV	Verify that the Distributed Services Time Out Value timers match exactly.
E_D_TOV	Verify that the Error Detect Time Out Value timers match exactly.
R_A_TOV	Verify that the Resource Allocation Time Out Value timers match exactly.
Trunking	Trunking is not supported between two different vendor's switches. This feature may be disabled on a per port or per switch basis.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Table 22-2 Changes in Switch Operation When Interoperability Is Enabled (continued)

Switch Feature	Changes if Interoperability Is Enabled
Default zone	The default zone operation 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 On a Brocade switch, use the cfgsave command to save fabric-wide zoning configuration. This command does not have any effect on Cisco Nexus 5000 Series switches if they are part of the same fabric. You must explicitly save the configuration on each Cisco Nexus 5000 Series switch.
Zone propagation	Some vendors do not pass the full zone configuration to other switches, only the active zone set gets passed. Verify that the active zone set or zone configuration has correctly propagated to the other switches in the fabric.
VSAN	Interop mode only affects the specified VSAN. Note Interop modes cannot be enabled on FICON-enabled VSANs.
TE ports and SAN port channels	TE ports and SAN port channels cannot be used to connect Cisco switches to non-Cisco SAN switches. Only E ports can be used to connect to non-Cisco SAN switches. TE ports and SAN port channels can still be used to connect a Cisco switch to other Cisco SAN 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. Cisco Nexus 5000 Series switches have the capability to restart only the domain manager process for the affected VSAN and not the entire switch.
Name server	Verify that all vendors have the correct values in their respective name server database.

Configuring Interop Mode 1

The interop mode1 in Cisco Nexus 5000 Series switches can be enabled disruptively or nondisruptively.



Note

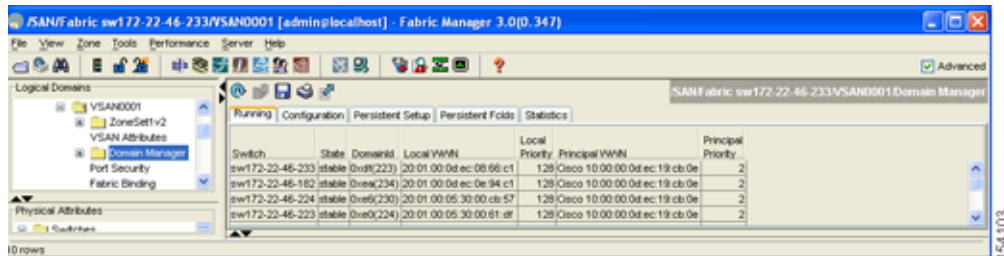
Brocade's **msplmgmtdeactivate** command must explicitly be run prior to connecting from a Brocade switch to either Cisco Nexus 5000 Series switches or to McData switches. This command uses Brocade proprietary frames to exchange platform information, which Cisco Nexus 5000 Series switches or McData switches do not understand. Rejecting these frames causes the common E ports to become isolated.

To configure interop mode 1 for a VSAN using Fabric Manager, perform this task:

Send comments to nx5000-docfeedback@cisco.com

-
- Step 1** Choose **VSANxxx > VSAN Attributes** from the Logical Domains pane.
- Step 2** Choose **Interop-1** from the Interop drop-down list.
- Step 3** Click **Apply Changes** to save this interop mode.
- Step 4** Expand **VSANxxx**, and then choose **Domain Manager** from the Logical Domains pane. You see the Domain Manager configuration in the Information pane as shown in [Figure 22-5](#).

Figure 22-5 Domain Manager Configuration



- Step 5** Set the Domain ID in the range of 97 (0x61) through 127 (0x7F).
- Click the **Configuration** tab.
 - Click in the Config Domain ID column under the Configuration tab.
 - Click the **Running** tab and verify that the change has been made.



Note The domain ID range limit is to accommodate McData switches.



Note When changing the domain ID, the FC IDs assigned to N ports also change.

- Step 6** Change the Fibre Channel timers (if they have been changed from the system defaults).



Note The Cisco, Brocade, and McData FC error detect (ED_TOV) and resource allocation (RA_TOV) timers default to the same values. They can be changed if needed. The RA_TOV default is 10 seconds, and the ED_TOV default is 2 seconds. Per the FC-SW2 standard, these values must be the same on each switch within the fabric.

- Expand **Switches > FC Services**, and then choose **Timers and Policies**. You see the timer settings in the Information pane.
 - Click **Change Timeouts** to modify the time-out values.
 - Click **Apply** to save the new time-out values.
- Step 7** (Optional) Choose **VSANxxx > Domain Manager**, click the **Configuration** tab, and choose **disruptive** or **nonDisruptive** in the Restart drop-down list to restart the domain.
-

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Verifying Interoperating Status

This section highlights the steps used to verify if the fabric is up and running in interoperability mode. To verify the interoperability status of the Cisco Cisco Nexus 5000 Series switch using Fabric Manager, perform this task:

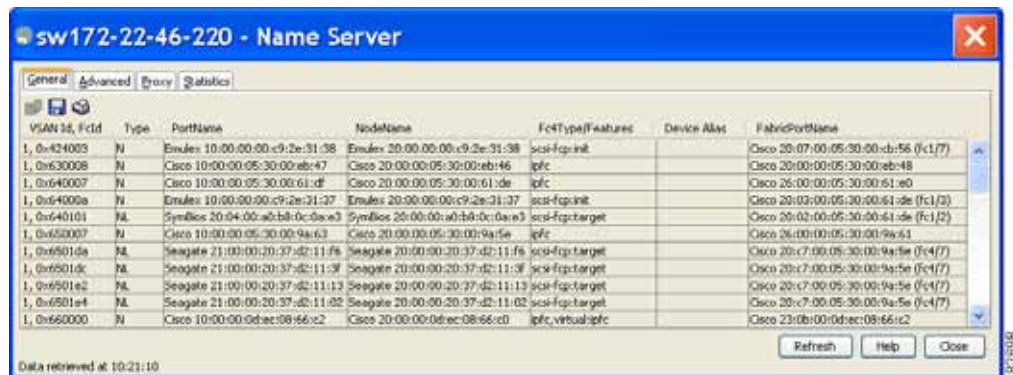
- Step 1** Choose **Switches** in the Physical Attributes pane and check the release number in the Information pane to verify the Cisco SAN-OS release.
- Step 2** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical** to verify the interface modes for each switch.
- Step 3** Expand **Fabricxx** in the Logical Domains pane, and then choose **All VSANs** to verify the interop mode for all VSANs.
- Step 4** Expand **Fabricxx > All VSANs**, and then choose **Domain Manager** to verify the domain IDs, local, and principal sWWNs for all VSANs (see [Figure 22-6](#)).

Figure 22-6 Domain Manager Information



- Step 5** Using Device Manager, choose **FC > Name Server** to verify the name server information. You see the Name Server dialog box as shown in [Figure 22-7](#).

Figure 22-7 Name Server Dialog Box



- Step 6** Click **Close** to close the dialog box.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)



Note

The Cisco switch name server shows both local and remote entries, and does not time out the entries.

Default Settings

Table 22-3 lists the default settings for the features included in this chapter.

Table 22-3 Default Settings for Advanced Features

Parameters	Default
CIM server	Disabled
CIM server security protocol	HTTP
D_S_TOV	5,000 milliseconds
E_D_TOV	2,000 milliseconds
R_A_TOV	10,000 milliseconds
Timeout period to invoke fctrace	5 seconds
Number of frame sent by the fcping feature	5 frames
Remote capture connection protocol	TCP
Remote capture connection mode	Passive
Local capture frame limits	10 frames
FC ID allocation mode	Auto mode
Loop monitoring	Disabled
Interop mode	Disabled



Configuring FC-SP and DHCHAP

Fibre Channel Security Protocol (FC-SP) capabilities provide switch-to-switch and host-to-switch authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) is an FC-SP protocol that provides authentication between Cisco Nexus 5000 Series switches and other devices. DHCHAP consists of the CHAP protocol combined with the Diffie-Hellman exchange.

This chapter includes the following sections:

- [Information About Fabric Authentication, page 23-1](#)
- [DHCHAP, page 23-2](#)
- [Default Settings, page 23-10](#)


Information About Fabric Authentication

All Cisco Nexus 5000 Series switches enable fabric-wide authentication from one switch to another switch, or from a switch to a host. These switch and host authentications are performed locally or remotely in each fabric. As storage islands are consolidated and migrated to enterprise-wide fabrics new security challenges arise. The approach of securing storage islands cannot always be guaranteed in enterprise-wide fabrics. For example, in a campus environment with geographically distributed switches, someone could maliciously interconnect incompatible switches or you could accidentally do so, resulting in Inter-Switch Link (ISL) isolation and link disruption.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Cisco Nexus 5000 Series switches support authentication features to address physical security (see [Figure 23-1](#)).

Figure 23-1 Switch and Host Authentication



106209

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, which prevents unauthorized devices from accessing the switch.



Note

The terms FC-SP and DHCHAP are used interchangeably in this chapter.

DHCHAP is a mandatory password-based, key-exchange authentication protocol that supports both switch-to-switch and host-to-switch authentication. DHCHAP negotiates hash algorithms and DH groups before performing authentication. It supports MD5 and SHA-1 algorithm-based authentication.

Send comments to nx5000-docfeedback@cisco.com

To configure DHCHAP authentication using the local password database, perform this task:

-
- | | |
|---------------|--|
| Step 1 | Enable DHCHAP. |
| Step 2 | Identify and configure the DHCHAP authentication modes. |
| Step 3 | Configure the hash algorithm and DH group. |
| Step 4 | Configure the DHCHAP password for the local switch and other switches in the fabric. |
| Step 5 | Configure the DHCHAP timeout value for reauthentication. |
| Step 6 | Verify the DHCHAP configuration. |
-

This section includes the following topics:

- [DHCHAP Compatibility with Fibre Channel Features, page 23-3](#)
- [About Enabling DHCHAP, page 23-4](#)
- [Enabling DHCHAP, page 23-4](#)
- [About DHCHAP Authentication Modes, page 23-4](#)
- [Configuring the DHCHAP Mode, page 23-5](#)
- [About the DHCHAP Hash Algorithm, page 23-6](#)
- [Configuring the DHCHAP Hash Algorithm, page 23-6](#)
- [About the DHCHAP Group Settings, page 23-6](#)
- [Configuring the DHCHAP Group Settings, page 23-6](#)
- [About the DHCHAP Password, page 23-7](#)
- [Configuring DHCHAP Passwords for the Local Switch, page 23-7](#)
- [About Password Configuration for Remote Devices, page 23-8](#)
- [Configuring DHCHAP Passwords for Remote Devices, page 23-8](#)
- [About the DHCHAP Timeout Value, page 23-8](#)
- [Configuring the DHCHAP Timeout Value, page 23-9](#)
- [Configuring DHCHAP AAA Authentication, page 23-9](#)
- [Enabling FC-SP on ISLs, page 23-9](#)

DHCHAP Compatibility with Fibre Channel Features

This section identifies the impact of configuring the DHCHAP feature along with existing Cisco NX-OS features:

- SAN port channel interfaces—If DHCHAP is enabled for ports belonging to a SAN port channel, DHCHAP authentication is performed at the physical interface level, not at the port channel level.
- Port security or fabric binding—Fabric-binding policies are enforced based on identities authenticated by DHCHAP.
- VSANs—DHCHAP authentication is not done on a per-VSAN basis.

Send comments to nx5000-docfeedback@cisco.com

About Enabling DHCHAP

By default, the DHCHAP feature is disabled in all Cisco Nexus 5000 Series switches.

You must explicitly enable the DHCHAP feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.


Enabling DHCHAP

To enable DHCHAP for a Cisco MDS switch using Fabric Manager, perform this task:

Step 1 Expand **Switches**, expand **Security**, and then choose **FC-SP**.

You see the FC-SP (DHCHAP) configuration in the Information pane as shown in [Figure 23-2](#).

Figure 23-2 FC-SP Configuration



Switch	Status	Command	LastCommand	Result
sw172-22-46-220	disabled	noSelection	noSelection	none
sw172-22-46-224	enabled	noSelection	noSelection	none
sw172-22-46-221	enabled	noSelection	noSelection	none
sw172-22-46-225	disabled	noSelection	noSelection	none
sw172-22-46-223	enabled	noSelection	noSelection	none
sw172-22-46-233	enabled	noSelection	noSelection	none
sw172-22-46-222	enabled	noSelection	noSelection	none
sw172-22-46-174	disabled	noSelection	noSelection	none

The **Control** tab is the default. You see the FC-SP enable state for all switches in the fabric.

Step 2 In the Command drop-down list, choose **enable** for all switches that you want to enable FC-SP on.

Step 3 Click the **Apply Changes** icon to enable FC-SP and DHCHAP on the selected switches.

About DHCHAP Authentication Modes

The DHCHAP authentication status for each interface depends on the configured DHCHAP port mode.

When the DHCHAP feature is enabled in a switch, each Fibre Channel interface or FCIP interface may be configured to be in one of four DHCHAP port modes:

- **On**—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the link is placed in an isolated state.
- **Auto-Active**—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software continues with the rest of the initialization sequence.
- **Auto-Passive (default)**—The switch does not initiate DHCHAP authentication, but participates in DHCHAP authentication if the connecting device initiates DHCHAP authentication.
- **Off**—The switch does not support DHCHAP authentication. Authentication messages sent to ports in this mode return error messages to the initiating switch.

Send comments to nx5000-docfeedback@cisco.com

**Note**

Whenever DHCHAP port mode is changed to a mode other than the Off mode, reauthentication is performed.

Table 23-1 identifies switch-to-switch authentication between two Cisco switches in various modes.

Table 23-1 DHCHAP Authentication Status Between Two MDS Switches

Switch N DHCHAP Modes	Switch 1 DHCHAP Modes			
	on	auto-active	auto-passive	off
on	FC-SP authentication is performed.	FC-SP authentication is performed.	FC-SP authentication is performed.	Link is brought down.
auto-Active			FC-SP authentication is <i>not</i> performed.	FC-SP authentication is <i>not</i> performed.
auto-Passive				
off	Link is brought down.	FC-SP authentication is <i>not</i> performed.		

Configuring the DHCHAP Mode

To configure the DHCHAP mode for a particular interface using Fabric Manager, perform this task:

- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **FC Physical**. You see the interface configuration in the Information Pane.
- Step 2** Click the FC-SP tab. You see the FC-SP (DHCHAP) configuration in the Information pane as shown in [Figure 23-3](#).

Figure 23-3 FC-SP (DHCHAP) Interface Modes

Switch	Interface	Mode	ReAuth Interval (hr)	ReAuth Start	Auth Successes	Auth Fails	Auth Bypasses
c-186	fc1/1	autoPassive	0	<input type="checkbox"/>	0	0	0
c-186	fc1/2	autoPassive	0	<input type="checkbox"/>	0	0	0
c-186	fc1/4	autoPassive	0	<input type="checkbox"/>	0	0	0

- Step 3** In the **Mode** drop-down list, choose **DHCHAP authentication mode** for each interface that you want to support FC-SP.
- Step 4** Click the **Apply Changes** icon to save these DHCHAP port mode settings.

Send comments to nx5000-docfeedback@cisco.com

About the DHCHAP Hash Algorithm

Cisco SAN switches support a default hash algorithm priority list of MD5 followed by SHA-1 for DHCHAP authentication.



Tip

If you change the hash algorithm configuration, then change it globally for all switches in the fabric.



Caution

RADIUS and TACACS+ protocols always use MD5 for CHAP authentication. Using SHA-1 as the hash algorithm may prevent RADIUS and TACACS+ usage, even if these AAA protocols are enabled for DHCHAP authentication.

Configuring the DHCHAP Hash Algorithm

To configure the hash algorithm using Fabric Manager, perform this task:

Step 1 Choose **Switches > Security**, and then choose **FC-SP**.

Step 2 Click the **General/Password** tab.

You see the DHCHAP general settings mode for each switch as shown in [Figure 23-4](#).

Figure 23-4 General/ Password Tab

Switch	Timeout (sec)	DH-CHAP HashList	DH-CHAP GroupList	GenericPassword
sw172-22-46-224	30	md5-sha1	null:1536:1024:1280:2048	*****
sw172-22-46-223	30	md5-sha1	null:1536:1024:1280:2048	*****
sw172-22-46-222	30	md5-sha1	null:1536:1024:1280:2048	*****
sw172-22-46-233	30	md5-sha1	null:1536:1024:1280:2048	*****
sw172-22-46-221	30	md5-sha1	null:1536:1024:1280:2048	*****

Step 3 Change the DHCHAP HashList for each switch in the fabric.

Step 4 Click the **Apply Changes** icon to save the updated hash algorithm priority list.

About the DHCHAP Group Settings

All Cisco Nexus 5000 Series switches support all DHCHAP groups specified in the standard: 0 (null DH group, which does not perform the Diffie-Hellman exchange), 1, 2, 3, or 4.



Tip

If you change the DH group configuration, change it globally for all switches in the fabric.

Configuring the DHCHAP Group Settings

To change the DH group settings using Fabric Manager, perform this task:

Send comments to nx5000-docfeedback@cisco.com

-
- Step 1** Expand **Switches > Security**, and then choose **FC-SP**.
- Step 2** Click the **General/Password** tab.
- Step 3** Change the DHCHAP GroupList for each switch in the fabric.
- Step 4** Click the **Apply Changes** icon to save the updated hash algorithm priority list.
-

About the DHCHAP Password

DHCHAP authentication in each direction requires a shared secret password between the connected devices. To do this, you can use one of three configurations to manage passwords for all switches in the fabric that participate in DHCHAP:

- Configuration 1—Use the same password for all switches in the fabric. This is the simplest configuration. When you add a new switch, you use the same password to authenticate that switch in this fabric. It is also the most vulnerable configuration if someone from the outside maliciously attempts to access any one switch in the fabric.
- Configuration 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.
- Configuration 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 configuration requires considerable password maintenance by the user.



Note All passwords are restricted to 64 alphanumeric characters and can be changed, but not deleted.



Tip We recommend using RADIUS or TACACS+ for fabrics with more than five switches. If you need to use a local password database, you can continue to do so using Configuration 3 and using the Cisco MDS 9000 Family Fabric Manager to manage the password database.

Configuring DHCHAP Passwords for the Local Switch

To configure the DHCHAP password for the local switch using Fabric Manager, perform this task:

-
- Step 1** Expand **Switches > Security**, and then choose **FC-SP**.
You see the FC-SP configuration in the Information pane.
- Step 2** Click the **Local Passwords** tab.
- Step 3** Click the **Create Row** icon to create a new local password.
You see the Create Local Passwords dialog box.
- Step 4** (Optional) Check the switches that you want to configure the same local password on.
- Step 5** Select the switch WNN and fill in the Password field.

Send comments to nx5000-docfeedback@cisco.com

Step 6 Click **Create** to save the updated password.

About Password Configuration for Remote Devices

You can configure passwords in the local authentication database for other devices in a fabric. The other devices are identified by their device name, which is also known as the switch WWN or device WWN. The password is restricted to 64 characters and can be specified in clear text (0) or in encrypted text (7).



Note

The switch WWN identifies the physical switch. This WWN is used to authenticate the switch and is different from the VSAN node WWN.

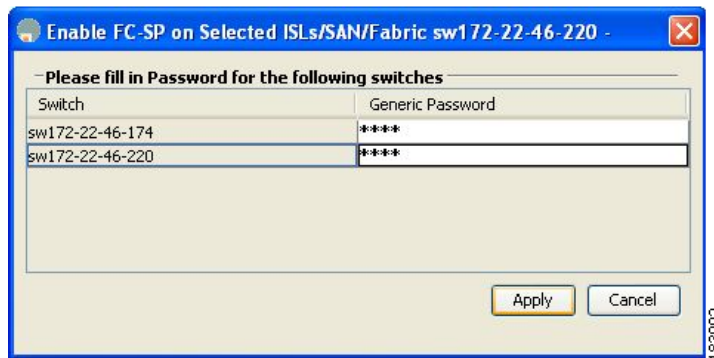
Configuring DHCHAP Passwords for Remote Devices

To locally configure the remote DHCHAP password for another switch in the fabric using Fabric Manager, perform this task:

Step 1 Right-click an ISL and choose **Enable FC-SP** from the drop-down list.

You see the Enable FC-SP dialog box as shown in [Figure 23-5](#).

Figure 23-5 Enable FC-SP Dialog Box



Step 2 Click **Apply** to save the updated password.

About the DHCHAP Timeout Value

During the DHCHAP protocol exchange, if the Cisco Nexus 5000 Series 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.

Send comments to nx5000-docfeedback@cisco.com

When changing the timeout value, consider the following factors:

- The existing RADIUS and TACACS+ timeout values.
- The same value must also be configured on all switches in the fabric.

Configuring the DHCHAP Timeout Value

To configure the DHCHAP timeout value using Fabric Manager, perform this task:

-
- Step 1** Expand **Switches > Security**, and then choose **FC-SP**.
You see the FC-SP configuration in the Information pane.
- Step 2** Click the **General/Password** tab.
You see the DHCHAP general settings mode for each switch as shown in [Figure 23-6](#).

Figure 23-6 General/Password Tab

Switch	Timeout (sec)	DH-CHAP HashList	DH-CHAP GroupList	GenericPassword
c-186	30	md5:sha1	null:1536:1024:1280:2048	
sw-189	30	md5:sha1	null:1536:1024:1280:2048	

- Step 3** Change the DHCHAP timeout value for each switch in the fabric.
- Step 4** Click the **Apply Changes** icon to save the updated information.
-

Configuring DHCHAP AAA Authentication

You can configure AAA authentication to use a RADIUS or TACACS+ server group. If AAA authentication is not configured, local authentication is used by default.

To configure the AAA authentication, see the *Cisco Cisco Nexus 5000 Series CLI Configuration Guide*.

Enabling FC-SP on ISLs

There is an ISL pop-up menu in Fabric Manager called Enable FC-SP that enables FC-SP on switches at either end of the ISL. You are prompted for an FC-SP generic password, then asked to set FC-SP interface mode to On for affected ports. Right-click an ISL and click **Enable FC-SP** to access this feature.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Default Settings

Table 23-2 lists the default settings for all fabric security features in any switch.

Table 23-2 *Default Fabric Security Settings*

Parameters	Default
DHCHAP feature	Disabled
DHCHAP hash algorithm	A priority list of MD5 followed by SHA-1 for DHCHAP authentication
DHCHAP authentication mode	Auto-passive
DHCHAP group default priority exchange order	0, 4, 1, 2, and 3, respectively
DHCHAP timeout value	30 seconds



Configuring Port Security

Cisco Nexus 5000 Series switches provide port security features that reject intrusion attempts and report these intrusions to the administrator.



Note

Port security is supported on virtual Fibre Channel ports and physical Fibre Channel ports.

This chapter includes the following sections:

- [Information About Port Security, page 24-1](#)
- [Configuring Port Security, page 24-3](#)
- [Enabling Port Security, page 24-5](#)
- [Port Security Activation, page 24-6](#)
- [Auto-Learning, page 24-10](#)
- [Port Security Manual Configuration, page 24-13](#)
- [Port Security Configuration Distribution, page 24-15](#)
- [Database Merge Guidelines, page 24-18](#)
- [Database Interaction, page 24-18](#)
- [Default Settings, page 24-21](#)

Information About Port Security

Typically, any Fibre Channel device in a SAN can attach to any SAN switch port and access SAN services based on zone membership. Port security features prevent unauthorized access to a switch port in the Cisco Nexus 5000 Series switch, using the following methods:

- Login requests from unauthorized Fibre Channel devices (N ports) and switches (xE ports) are rejected.
- All intrusion attempts are reported to the SAN administrator through system messages.
- Configuration distribution uses the CFS infrastructure, and is limited to those switches that are CFS capable. Distribution is disabled by default.
- Configuring the port security policy requires the Storage Protocol Services license.

Send comments to nx5000-docfeedback@cisco.com

This section includes the following topics:

- [Port Security Enforcement, page 24-2](#)
- [About Auto-Learning, page 24-2](#)
- [Port Security Activation, page 24-3](#)

Port Security Enforcement

To enforce port security, configure the devices and switch port interfaces through which each device or switch is connected, and activate the configuration.

- Use the port world wide name (pWWN) or the node world wide name (nWWN) to specify the N port connection for each device.
- Use the switch world wide name (sWWN) to specify the xE port connection for each switch.

Each N and xE port can be configured to restrict a single port or a range of ports.

Enforcement of port security policies are done on every activation and when the port tries to come up.

The port security feature uses two databases to accept and implement configuration changes.

- Configuration database—All configuration changes are stored in the configuration database.
- Active database—The database currently enforced by the fabric. The port security feature requires all devices connecting to a switch to be part of the port security active database. The software uses this active database to enforce authorization.

About Auto-Learning

You can instruct the switch to automatically learn (auto-learn) the port security configurations over a specified period. This feature allows any Cisco Nexus 5000 Series switch to automatically learn about devices and switches that connect to it. Use this feature when you activate the port security feature for the first time as it saves tedious manual configuration for each port. You must configure auto-learning on a per-VSAN basis. If enabled, devices and switches that are allowed to connect to the switch are automatically learned, even if you have not configured any port access.

When auto-learning is enabled, learning happens only for the devices or interfaces that were not already logged into the switch. Learned entries on a port are cleaned up after you shut down that port if auto-learning is still enabled.

Learning does not override the existing configured port security policies. For example, if an interface is configured to allow a specific pWWN, then auto-learning will not add a new entry to allow any other pWWN on that interface. All other pWWNs will be blocked even in auto-learning mode.

No entries are learned for a port in the shutdown state.

When you activate the port security feature, auto-learning is also automatically enabled.



Note

If you enable auto-learning before activating port security, you cannot activate port security until auto-learning is disabled.

Send comments to nx5000-docfeedback@cisco.com

Port Security Activation

By default, the port security feature is not activated in Cisco Nexus 5000 Series switches.

When you activate the port security feature, the following operations occur:

- Auto-learning is also automatically enabled, which means:
 - From this point, auto-learning happens only for the devices or interfaces that were not logged into the switch.
 - You cannot activate the database until you disable auto-learning.
- All the devices that are already logged in are learned and are added to the active database.
- All entries in the configured database are copied to the active database.

After the database is activated, subsequent device login is subject to the activated port bound WWN pairs, excluding the auto-learned entries. You must disable auto-learning before the auto-learned entries become activated.

When you activate the port security feature, auto-learning is also automatically enabled. You can choose to activate the port security feature and disable auto-learning.



Tip

If a port is shut down because of a denied login attempt, and you subsequently configure the database to allow that login, the port does not come up automatically. You must explicitly enter a **no shutdown** CLI command to bring that port back online.

Configuring Port Security

The steps to configure port security depend on which features you are using. Auto-learning works differently if you are using CFS distribution.

This section includes the following topics:

- [Configuring Port Security with Auto-Learning and CFS Distribution, page 24-3](#)
- [Configuring Port Security with Auto-Learning without CFS, page 24-4](#)
- [Configuring Port Security with Manual Database Configuration, page 24-5](#)

Configuring Port Security with Auto-Learning and CFS Distribution

To configure port security, using auto-learning and CFS distribution, perform this task:

-
- | | |
|---------------|---|
| Step 1 | Enable port security.
See the “Enabling Port Security” section on page 24-5. |
| Step 2 | Enable CFS distribution.
See the “Enabling Distribution” section on page 24-16. |
| Step 3 | Activate port security on each VSAN.
This action turns on auto-learning by default. See the “Activating Port Security” section on page 24-7. |
| Step 4 | Issue a CFS commit to copy this configuration to all switches in the fabric. |

Send comments to nx5000-docfeedback@cisco.com

See the “[Committing the Changes](#)” section on page 24-17. All switches have port security activated with auto-learning enabled.

Step 5 Wait until all switches and all hosts are automatically learned.

Step 6 Disable auto-learn on each VSAN.

See the “[Disabling Auto-Learning](#)” section on page 24-11.

Step 7 Issue a CFS commit to copy this configuration to all switches in the fabric.

See the “[Committing the Changes](#)” section on page 24-17. The auto-learned entries from every switch are combined into a static active database that is distributed to all switches.

Step 8 Copy the active database to the configure database on each VSAN.

See the “[Copying the Port Security Database](#)” section on page 24-20.

Step 9 Issue a CFS commit to copy this configuration to all switches in the fabric.

See the “[Committing the Changes](#)” section on page 24-17. This ensures that the configure database is the same on all switches in the fabric.

Step 10 Copy the running configuration to the startup configuration, using the fabric option.

This step saves the port security configure database to the startup configuration on all switches in the fabric.

Configuring Port Security with Auto-Learning without CFS

To configure port security using auto-learning without CFS, perform this task:

Step 1 Enable port security.

See the “[Enabling Port Security](#)” section on page 24-5.

Step 2 Activate port security on each VSAN, which turns on auto-learning by default.

See the “[Activating Port Security](#)” section on page 24-7.

Step 3 Wait until all switches and all hosts are automatically learned.

Step 4 Disable auto-learn on each VSAN.

See the “[Disabling Auto-Learning](#)” section on page 24-11.

Step 5 Copy the active database to the configure database on each VSAN.

See the “[Copying the Port Security Database](#)” section on page 24-20.

Step 6 Copy the running configuration to the startup configuration, which saves the port security configuration database to the startup configuration.

Step 7 Repeat [Step 1](#) through [Step 6](#) for all switches in the fabric.

Send comments to nx5000-docfeedback@cisco.com

Configuring Port Security with Manual Database Configuration

To configure port security and manually configure the port security database, perform this task:

-
- Step 1** Enable port security.
See the “[Enabling Port Security](#)” section on page 24-5.
 - Step 2** Manually configure all port security entries into the configure database on each VSAN.
See the “[Configuring Port Security with Manual Database Configuration](#)” section on page 24-5.
 - Step 3** Activate port security on each VSAN. This turns on auto-learning by default.
See the “[Disabling Auto-Learning](#)” section on page 24-11.
 - Step 4** Disable auto-learn on each VSAN.
See the “[Disabling Auto-Learning](#)” section on page 24-11.
 - Step 5** Copy the running configuration to the startup configuration, which saves the port security configuration database to the startup configuration.
 - Step 6** Repeat [Step 1](#) through [Step 5](#) for all switches in the fabric.
-

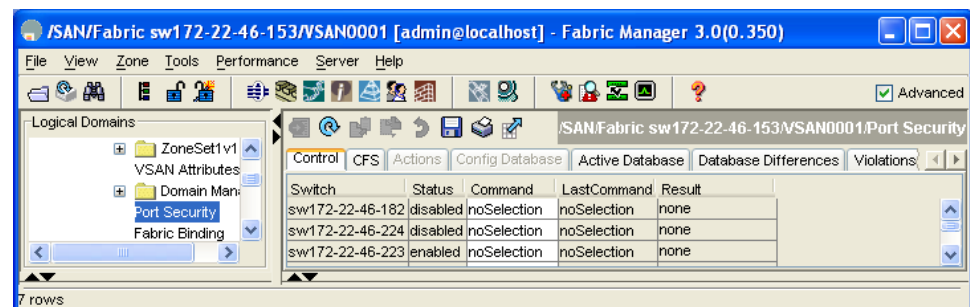
Enabling Port Security

By default, the port security feature is disabled in Cisco Nexus 5000 Series switches.

To enable port security using Fabric Manager, perform this task:

-
- Step 1** Expand a VSAN, and then choose **Port Security** in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane as shown in [Figure 24-1](#).

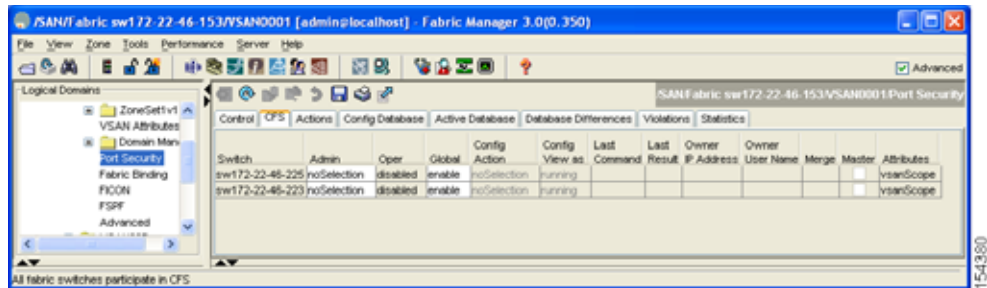
Figure 24-1 Port Security Configuration



- Step 2** Click the **CFS** tab.
You see the information shown in [Figure 24-2](#).

Send comments to nx5000-docfeedback@cisco.com

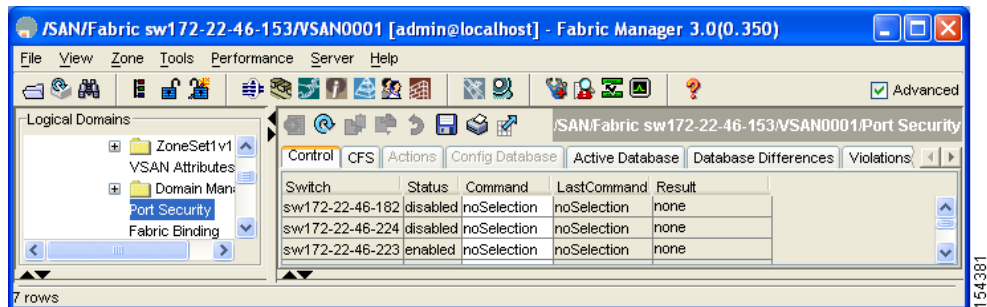
Figure 24-2 Port Security CFS



- Step 3** Enable CFS on all participating switches in the VSAN by clicking each entry in the Global column and selecting **enable**.
- Step 4** Click **Apply Changes** to enable CFS distribution for the port security feature.
- Step 5** Click the **Control** tab.

You see the port security enable state for all switches in the selected VSAN as shown in [Figure 24-3](#).

Figure 24-3 Port Security Configuration



- Step 6** In the Status column, choose **enable** for each switch in the VSAN.
- Step 7** Click the **CFS** tab and in the Command column choose **commit** on all participating switches in the VSAN.
- Step 8** Click **Apply Changes** to distribute the enabled port security to all switches in the VSAN.

Port Security Activation

This section includes the following topics:

- [Activating Port Security, page 24-7](#)
- [Database Activation Rejection, page 24-7](#)
- [Forcing Port Security Activation, page 24-8](#)
- [Database Reactivation, page 24-8](#)
- [Copying an Active Database to the Config Database, page 24-9](#)
- [Displaying Activated Port Security Settings, page 24-9](#)

Send comments to nx5000-docfeedback@cisco.com

- [Displaying Port Security Statistics, page 24-9](#)
- [Displaying Port Security Violations, page 24-10](#)

Activating Port Security

To activate port security using Fabric Manager, perform this task:

-
- Step 1** Expand a **VSAN** and then choose **Port Security** in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the **Actions** tab.
- Step 3** In the Action column under Activation, choose the switch or VSAN on which you want to activate port security. You see a drop-down list with the following options:
- **activate**—Valid port security settings are activated.
 - **activate (TurnLearningOff)**—Valid port security settings are activated and auto-learn turned off.
 - **forceActivate**—Activation is forced.
 - **forceActivate(TurnLearningOff)**—Activation is forced and auto-learn is turned off.
 - **deactivate**—All currently active port security settings are deactivated.
 - **NoSelection**— No action is taken.
- Step 4** Set the Action field you want for that switch.
- Step 5** Uncheck the **AutoLearn** check box for each switch in the VSAN to disable auto-learning.
- Step 6** Click the **CFS** tab and set the command column to **commit** on all participating switches in the VSAN.
- Step 7** Click **Apply Changes** in Fabric Manager to save these changes.
-

Database Activation Rejection

Database activation is rejected in the following cases:

- Missing or conflicting entries exist in the configuration database but not in the active database.
- The auto-learning feature was enabled before the activation. To reactivate a database in this state, disable auto-learning.
- The exact security is not configured for each port channel member.
- The configured database is empty but the active database is not.

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed by forcing the port security activation.

Send comments to nx5000-docfeedback@cisco.com

Forcing Port Security Activation

If the port security activation request is rejected, you can force the activation.



Note

If you force the activation, existing devices are logged out if they violate the active database.

To forcefully activate the port security database using Fabric Manager, perform this task:

-
- Step 1** Expand a **VSAN** and then choose **Port Security** in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
 - Step 2** Click the **Actions** tab.
 - Step 3** In the **Action** column under Activation, choose the switch or VSAN on which you want to activate port security and choose the **forceactivate** option.
 - Step 4** Choose the Action field you want for that switch.
 - Step 5** Click the **CFS** tab and in the command column choose **commit** for all participating switches in the VSAN.
 - Step 6** Click **Apply Changes** in Fabric Manager to save these changes.
-

Database Reactivation



Tip

If auto-learning is enabled, and you cannot activate the database, you will not be allowed to proceed.

To reactivate the port security database using Fabric Manager, perform this task:

-
- Step 1** Disable auto-learning.
 - Step 2** Copy the active database to the configured database.



Tip

If the active database is empty, you cannot perform this step.

- Step 3** Make the required changes to the configuration database.
 - Step 4** Activate the database.
-

Send comments to nx5000-docfeedback@cisco.com

Copying an Active Database to the Config Database

To copy the active database to the config database using Fabric Manager, perform this task:

-
- Step 1** Expand a **VSAN** and then choose **Port Security** in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
 - Step 2** Click the **Actions** tab.
You see the switches for that VSAN.
 - Step 3** Check the **CopyActive ToConfig** check box next to the switch for which you want to copy the database.
The active database is copied to the config database when the security setting is activated.
 - Step 4** Uncheck the **CopyActive ToConfig** check box if you do not want the database copied when the security setting is activated.
 - Step 5** Click the **CFS** tab and in the command column choose **commit** for all participating switches in the VSAN.
 - Step 6** Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.
-

Displaying Activated Port Security Settings

To display active port security settings using Fabric Manager, perform this task:

-
- Step 1** Expand a **VSAN** and then choose **Port Security** in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
 - Step 2** Click the **Active Database** tab.
You see the active port security settings for that VSAN.
-

Displaying Port Security Statistics

To display port security statistics using Fabric Manager, perform this task:

-
- Step 1** Expand a **VSAN** and then choose **Port Security** in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
 - Step 2** Click the **Statistics** tab.
You see the port security statistics for that VSAN.
-

Send comments to nx5000-docfeedback@cisco.com

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 this task:

-
- Step 1** Expand a **VSAN** and then choose **Port Security** in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the **Violations** tab.
You see the port security violations for that VSAN.
-

Auto-Learning

This section includes the following topics:

- [About Enabling Auto-Learning, page 24-10](#)
- [Enabling Auto-Learning, page 24-11](#)
- [Disabling Auto-Learning, page 24-11](#)
- [Auto-Learning Device Authorization, page 24-12](#)
- [Authorization Scenario, page 24-12](#)

About Enabling Auto-Learning

The state of the auto-learning configuration depends on the state of the port security feature:

- If the port security feature is not activated, auto-learning is disabled by default.
- If the port security feature is activated, auto-learning is enabled by default (unless you explicitly disabled this option).



Tip

If auto-learning is enabled on a VSAN, you can only activate the database for that VSAN by using the **force** option.

Send comments to nx5000-docfeedback@cisco.com

Enabling Auto-Learning

To enable auto-learning using Fabric Manager, perform this task:

- Step 1** Expand a **VSAN** and then choose **Port Security** in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane as shown in [Figure 24-4](#).

Figure 24-4 Port Security Configuration

Master	Action	Enabled	Result	LastChange	CopyActive ToConfig	AutoLearn	Clear Autolearned	AutoLearned Inters
SW172-22-46-220	NoSelection	False	success	n/a	<input type="checkbox"/>	<input checked="" type="checkbox"/>	NoSelection	

- Step 2** Click the **Actions** tab.
- Step 3** In the Action column under Activation, choose the switch or VSAN on which you want to activate port security. You see a drop-down list with the following options:
- **activate**—Valid port security settings are activated.
 - **activate (TurnLearningOff)**—Valid port security settings are activated and auto-learn turned off.
 - **forceActivate**—Activation is forced.
 - **forceActivate(TurnLearningOff)**—Activation is forced and auto-learn is turned off.
 - **deactivate**—All currently active port security settings are deactivated.
 - **NoSelection**— No action is taken.
- Step 4** Choose one of the port security options for that switch.
- Step 5** Check the **AutoLearn** check box for each switch in the VSAN to enable auto-learning.
- Step 6** Click the **Apply Changes** icon to save these changes.

Disabling Auto-Learning

To disable auto-learning using Fabric Manager, perform this task:

- Step 1** Expand a **VSAN** and then choose **Port Security** in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane (see [Figure 24-4](#)).
- Step 2** Click the **Actions** tab.
You see the switches for that VSAN.
- Step 3** Uncheck the **AutoLearn** check box next to the switch if you want to disable auto-learning.
- Step 4** Click the **Apply Changes** icon to save these changes.

Send comments to nx5000-docfeedback@cisco.com

Auto-Learning Device Authorization

Table 24-1 summarizes the authorized connection conditions for device requests.

Table 24-1 Authorized Auto-Learning Device Requests

Condition	Device (pWWN, nWWN, sWWN)	Requests Connection to	Authorization
1	Configured with one or more switch ports	A configured switch port	Permitted
2		Any other switch port	Denied
3	Not configured	A switch port that is not configured	Permitted if auto-learning enabled
4			Denied if auto-learning disabled
5	Configured or not configured	A switch port that allows any device	Permitted
6	Configured to log in to any switch port	Any port on the switch	Permitted
7	Not configured	A port configured with some other device	Denied

Authorization Scenario

Assume that the port security feature is activated and the following conditions are specified in the active database:

- A pWWN (P1) is allowed access through interface fc2/1 (F1).
- A pWWN (P2) is allowed access through interface fc2/2 (F1).
- A nWWN (N1) is allowed access through interface fc2/2 (F2).
- Any WWN is allowed access through interface vfc3/1 (F3).
- A nWWN (N3) is allowed access through any interface.
- A pWWN (P3) is allowed access through interface fc2/4 (F4).
- A sWWN (S1) is allowed access through interface fc3/1-3 (F10 to F13).
- A pWWN (P10) is allowed access through interface vfc4/1 (F11).

Table 24-2 summarizes the port security authorization results for this active database. The conditions listed refer to the conditions from Table 24-1.

Table 24-2 Authorization Results for Scenario

Device Connection Request	Authorization	Condition	Reason
P1, N2, F1	Permitted	1	No conflict.
P2, N2, F1	Permitted	1	No conflict.
P3, N2, F1	Denied	2	F1 is bound to P1/P2.
P1, N3, F1	Permitted	6	Wildcard match for N3.

Send comments to nx5000-docfeedback@cisco.com

Table 24-2 Authorization Results for Scenario (continued)

Device Connection Request	Authorization	Condition	Reason
P1, N1, F3	Permitted	5	Wildcard match for F3.
P1, N4, F5	Denied	2	P1 is bound to F1.
P5, N1, F5	Denied	2	N1 is only allowed on F2.
P3, N3, F4	Permitted	1	No conflict.
S1, F10	Permitted	1	No conflict.
S2, F11	Denied	7	P10 is bound to F11.
P4, N4, F5 (auto-learning on)	Permitted	3	No conflict.
P4, N4, F5 (auto-learning off)	Denied	4	No match.
S3, F5 (auto-learning on)	Permitted	3	No conflict.
S3, F5 (auto-learning off)	Denied	4	No match.
P1, N1, F6 (auto-learning on)	Denied	2	P1 is bound to F1.
P5, N5, F1 (auto-learning on)	Denied	7	Only P1 and P2 bound to F1.
S3, F4 (auto-learning on)	Denied	7	P3 paired with F4.
S1, F3 (auto-learning on)	Permitted	5	No conflict.
P5, N3, F3	Permitted	6	Wildcard (*) match for F3 and N3.
P7, N3, F9	Permitted	6	Wildcard (*) match for N3.

Port Security Manual Configuration

To configure port security on a Cisco Nexus 5000 Series switch, perform this task:

-
- Step 1** Identify the WWN of the ports that need to be secured.
See the [“Adding Authorized Port Pairs”](#) section on page 24-14.
- Step 2** Secure the fWWN to an authorized nWWN or pWWN.
- Step 3** Activate the port security database.
- Step 4** Verify your configuration.
-

This section includes the following topics:

- [WWN Identification Guidelines](#), page 24-14
- [Adding Authorized Port Pairs](#), page 24-14
- [Deleting Port Security Setting](#), page 24-15

Send comments to nx5000-docfeedback@cisco.com

WWN Identification Guidelines

If you decide to manually configure port security, note the following guidelines:

- Identify switch ports by the interface or by the fWWN.
- Identify devices by the pWWN or by the nWWN.
- If an N port is allowed to log in to SAN switch port F, then that N port can only log in through the specified F port.
- If an N port's nWWN is bound to an F port WWN, then all pWWNs in the N port are implicitly paired with the F port.
- TE port checking is done on each VSAN in the allowed VSAN list of the VSAN trunk port.
- All port channel xE ports must be configured with the same set of WWNs in the same SAN port channel.
- 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 configuration database can be modified without any effect on the active database.
- By saving the running configuration, you save the configuration database and activated entries in the active database. Learned entries in the active database are not saved.

Adding Authorized Port Pairs

After identifying the WWN pairs that need to be bound, add those pairs to the port security database.

**Tip**

Remote switch binding can be specified at the local switch. To specify the remote interfaces, you can use either the fWWN or sWWN-interface combination.

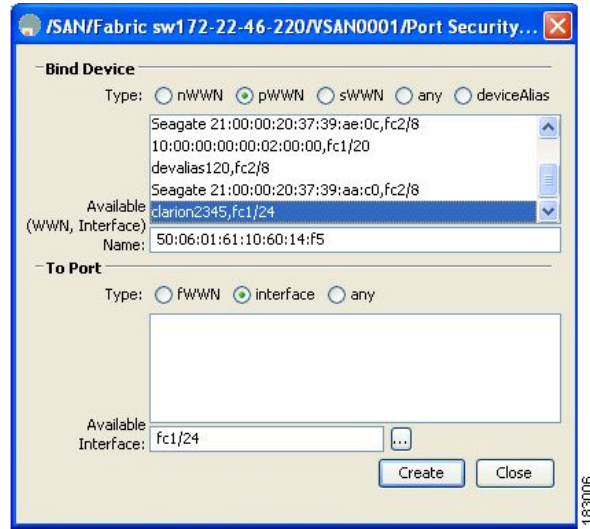
To add authorized port pairs for port security using Fabric Manager, perform this task:

- Step 1** Expand a **VSAN** and then choose **Port Security** in the Logical Domains pane.
- Step 2** Click the **Config Database** tab.
- Step 3** Click **Create Row** to add an authorized port pair.

You see the Create Port Security dialog box as shown in [Figure 24-5](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 24-5 Create Port Security Dialog Box



- Step 4** Double-click the device from the available list for which you want to create the port security setting.
- Step 5** Double-click the port from the available list to which you want to bind the device.
- Step 6** Click **Create** to create the port security setting.
- Step 7** Click the **Apply Changes** icon to save these changes.

Deleting Port Security Setting

To delete a port security setting from the configured database on a switch, perform this task:

- Step 1** Expand a **VSAN** and then choose **Port Security** in the Logical Domains pane.
- Step 2** Click the **Config Database** tab.
You see the configured port security settings for that VSAN.
- Step 3** Click in the row you want to delete.
- Step 4** Choose **Delete Row**.
You see the confirmation dialog box.
- Step 5** Click **Yes** to delete the row, or click **No** to close the confirmation dialog box without deleting the row.
- Step 6** Click the **Apply Changes** icon to save these changes.

Port Security Configuration Distribution

The port security feature uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management, provide a single point of configuration for the entire fabric in the VSAN, and enforce the port security policies throughout the fabric (see [Chapter 7, “Using Cisco Fabric Services”](#)).

Send comments to nx5000-docfeedback@cisco.com

This section contains the following topics:

- [Enabling Distribution, page 24-16](#)
- [Locking the Fabric, page 24-16](#)
- [Committing the Changes, page 24-17](#)
- [Activation and Auto-Learning Configuration Distribution, page 24-17](#)

Enabling Distribution

All the configurations performed in distributed mode are stored in a pending (temporary) database. If you modify the configuration, you need to commit or discard the pending database changes to the configurations. The fabric remains locked during this period. Changes to the pending database are not reflected in the configurations until you commit the changes.



Note

Port activation or deactivation and auto-learning enable or disable do not take effect until after a CFS commit if CFS distribution is enabled. Always follow any one of these operations with a CFS commit to ensure proper configuration. See the [“Activation and Auto-Learning Configuration Distribution” section on page 24-17](#).



Tip

We recommend that you perform a commit after you activate port security and after you enable auto learning.

To enable distribution using Fabric Manager, perform this task:

- Step 1** Expand a **VSAN** and then choose **Port Security** in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane as shown in [Figure 24-4](#).
- Step 2** Click the **Control** tab.
You see the switches for that VSAN.
- Step 3** In the **Command** column, choose **enable** or **disable** from the drop-down list.
- Step 4** Click the **Apply Changes** icon to save the changes.

Locking the Fabric

The first action that modifies the existing configuration creates the pending database and locks the feature in the VSAN. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database.

Send comments to nx5000-docfeedback@cisco.com

Committing the Changes

If you commit the changes made to the configurations, the configurations in the pending database are distributed to other switches. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

Activation and Auto-Learning Configuration Distribution

Activation and auto-learning configurations in distributed mode are remembered as actions to be performed when you commit the changes in the pending database.

Learned entries are temporary and do not have any role in determining if a login is authorized or not. As such, learned entries do not participate in distribution. When you disable learning and commit the changes in the pending database, the learned entries become static entries in the active database and are distributed to all switches in the fabric. After the commit, the active database on all switches are identical and learning can be disabled.

If the pending database contains more than one activation and auto-learning configuration when you commit the changes, the activation and auto-learning changes are consolidated and the resulting operation may change (see [Table 24-3](#)).

Table 24-3 Scenarios for Activation and Auto-learning Configurations in Distributed Mode

Scenario	Actions	Distribution = OFF	Distribution = ON
A and B exist in the configuration database, activation is not done and devices C,D are logged in.	1. You activate the port security database and enable auto-learning.	configuration database = {A,B} active database = {A,B, C ¹ , D*}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled}
	2. A new entry E is added to the configuration database.	configuration database = {A,B, E} active database = {A,B, C*, D*}	configuration database = {A,B} active database = {null} pending database = {A,B, E + activation to be enabled}
	3. You issue a commit.	Not applicable	configuration database = {A,B, E} active database = {A,B, E, C*, D*} pending database = empty

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Table 24-3 Scenarios for Activation and Auto-learning Configurations in Distributed Mode (continued)

Scenario	Actions	Distribution = OFF	Distribution = ON
A and B exist in the configuration database, activation is not done and devices C,D are logged in.	1. You activate the port security database and enable auto-learning.	configuration database = {A,B} active database = {A,B, C*, D*}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled}
	2. You disable learning.	configuration database = {A,B} active database = {A,B, C, D}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled + learning to be disabled}
	3. You issue a commit.	Not applicable	configuration database = {A,B} active database = {A,B} and devices C and D are logged out. This is equal to an activation with auto-learning disabled. pending database = empty

1. The * (asterisk) indicates learned entries.



Tip

In this case, we recommend that you perform a commit at the end of each operation: after you activate port security and after you enable auto-learning.

Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the active database. See the “[CFS Merge Support](#)” section on page 7-6 for detailed concepts.

When merging the database between two fabrics, follow these guidelines:

- Verify that the activation status and the auto-learning status is the same in both fabrics.
- Verify that the combined number of configurations for each VSAN in both databases does not exceed 2000.



Caution

If you do not follow these two conditions, the merge will fail. The next distribution will forcefully synchronize the databases and the activation states in the fabric.

Database Interaction

This section includes the following topics:

- [Database Scenarios, page 24-19](#)
- [Copying the Port Security Database, page 24-20](#)

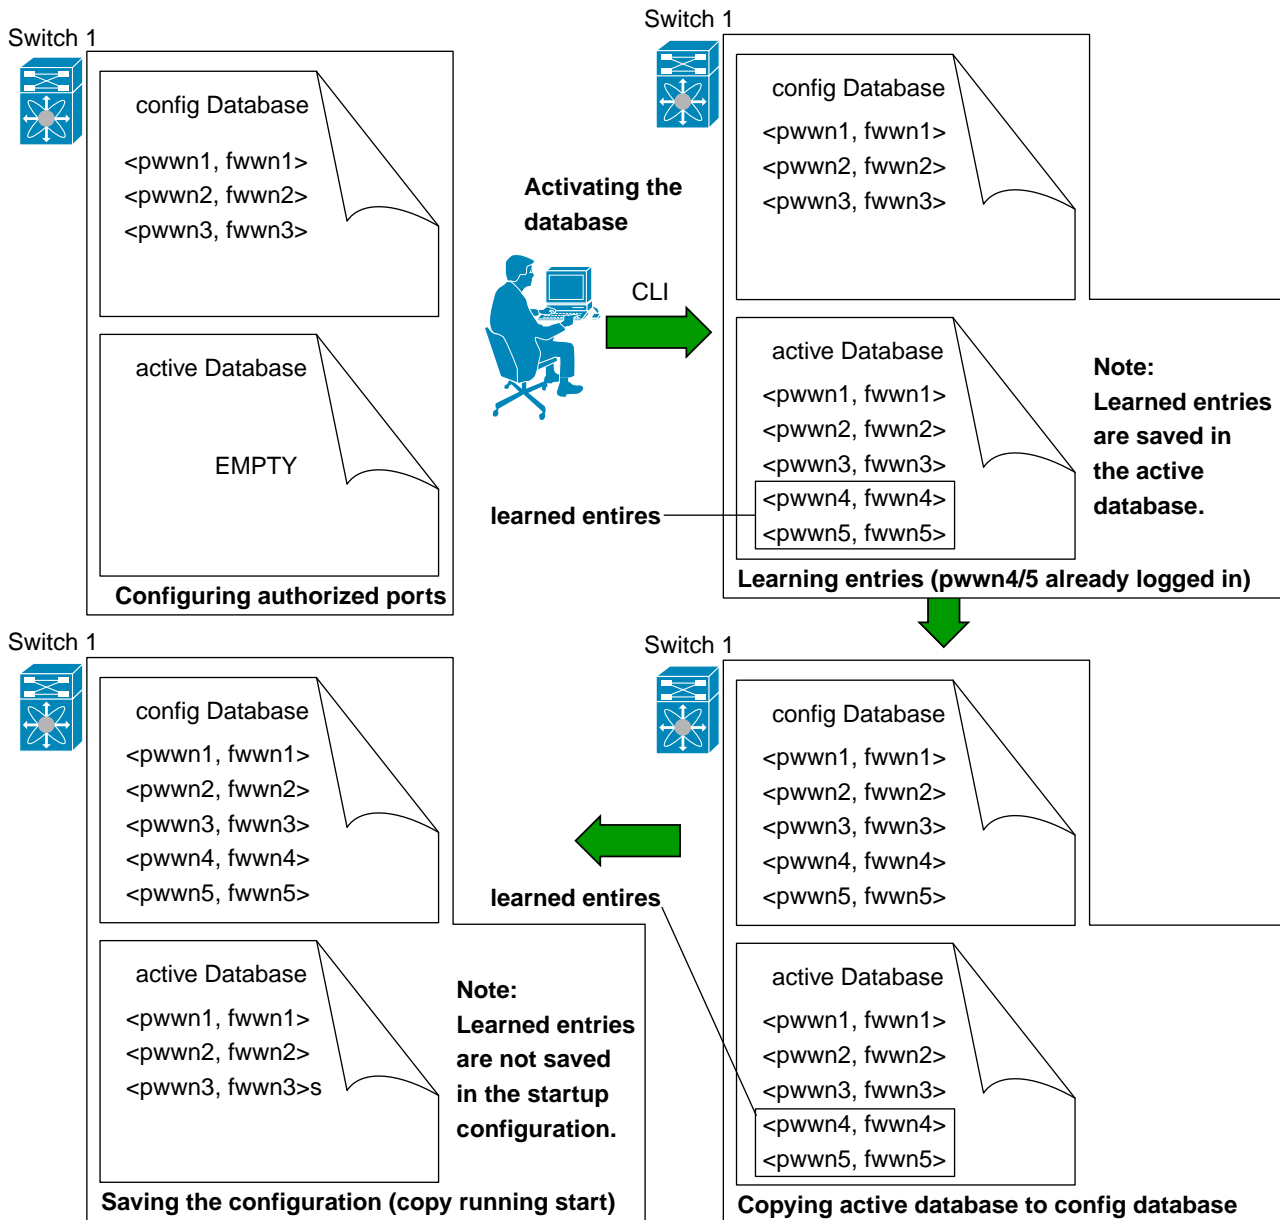
Send comments to nx5000-docfeedback@cisco.com

- [Deleting the Port Security Database, page 24-20](#)
- [Clearing the Port Security Database, page 24-21](#)

Database Scenarios

Figure 24-6 illustrates various scenarios showing the active database and the configuration database status based on port security configurations.

Figure 24-6 Port Security Database Scenarios



99301

Send comments to nx5000-docfeedback@cisco.com

Copying the Port Security Database



Tip

We recommend that you copy the active database to the config database after disabling auto-learning. This action will ensure that the configuration database is in synchronization with the active database. If distribution is enabled, this command creates a temporary copy (and consequently a fabric lock) of the configuration database. If you lock the fabric, you need to commit the changes to the configuration databases in all the switches.

To copy the active database to the configuration database, using Fabric Manager, perform this task:

-
- Step 1** Expand a **Fabric**, expand a **VSAN**, and then choose **Port Security** in the Logical Domains pane.
 - Step 2** Click the **Actions** tab.
You see all the configuration databases.
 - Step 3** Choose the appropriate configuration database and check the **Copy Active to Config** checkbox.
 - Step 4** Click the **Apply Changes** icon to save your changes.
-

To view the differences between the active database and the configuration database using Fabric Manager, perform this task:

-
- Step 1** Expand a **Fabric**, expand a **VSAN**, and then choose **Port Security** in the Logical Domains pane.
You see the Port Security information in the Information pane.
 - Step 2** Click the **Database Differences** tab.
You see all the configuration databases.
 - Step 3** Choose the appropriate configuration database. Choose the **Active** or **Config** option to compare the differences between the selected database and the active or configuration database.
 - Step 4** Click the **Apply Changes** icon to save your changes.
-

Deleting the Port Security Database



Tip

If the distribution is enabled, the deletion creates a copy of the database. An explicit deletion is required to actually delete the database.

To delete a port security database using Fabric Manager, perform this task:

-
- Step 1** Expand a **Fabric**, expand a **VSAN**, and then choose **Port Security** in the Logical Domains pane.
You see the Port Security information in the Information pane.
 - Step 2** Click the **Config Database** tab.
You see all the configuration databases.

Send comments to nx5000-docfeedback@cisco.com

- Step 3** Choose the appropriate configuration database and click the **Delete Row** button.
- Step 4** Click **Yes** if you want to delete the configuration database.

Clearing the Port Security Database

To clear all existing statistics from the port security database for a specified VSAN using Fabric Manager, perform this task:

- Step 1** Expand a **Fabric**, expand a **VSAN**, and then choose **Port Security** in the Logical Domains pane. You see the Port Security information in the Information pane (see [Figure 24-4](#)).
- Step 2** Click the **Statistics** tab. You see all the configuration databases.
- Step 3** Choose the appropriate configuration database and check the **Clear** option.
- Step 4** Click the **Apply Changes** icon to save your changes.

To clear any learned entries in the active database for a specified interface within a VSAN using Fabric Manager, perform this task:

- Step 1** Expand a **Fabric**, expand a **VSAN**, and then choose **Port Security** in the Logical Domains pane. You see the Port Security information in the Information pane.
- Step 2** Click the **Actions** tab. You see all the configuration databases.
- Step 3** Choose the appropriate configuration database and check the **AutoLearn** option.
- Step 4** Click the **Apply Changes** icon to save your changes.



Note You can clear the Statistics and the AutoLearn option only for switches that are local and do not acquire locks. Also, learned entries are only local to the switch and do not participate in distribution.

Default Settings

[Table 24-4](#) lists the default settings for all port security features in any switch.

Table 24-4 Default Security Settings

Parameters	Default
Auto-learn	Enabled if port security is enabled.

Send comments to nx5000-docfeedback@cisco.com

Table 24-4 Default Security Settings (continued)

Parameters	Default
Port security	Disabled.
Distribution	Disabled.
	Note Enabling distribution enables it on all VSANs in the switch.



Configuring Fabric Binding

This chapter describes the fabric binding feature provided in Cisco Nexus 5000 Series switches. It includes the following sections:

- [Information About Fabric Binding, page 25-1](#)
- [Configuring Fabric Binding, page 25-3](#)
- [Default Settings, page 25-10](#)

Information About Fabric Binding

The fabric binding feature ensures that ISLs are only enabled between specified switches in the fabric. Fabric binding is configured on a per-VSAN basis.

This feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. It uses the Exchange Fabric Membership Data (EFMD) protocol to ensure that the list of authorized switches is identical in all switches in the fabric.

This section includes the following topics:

- [Licensing Requirements, page 25-1](#)
- [Port Security Versus Fabric Binding, page 25-2](#)
- [Fabric Binding Enforcement, page 25-2](#)

Licensing Requirements

Fabric Binding requires the Storage Protocol Services license. For additional information, refer to the *Nexus 5000 Series Switch CLI Software Configuration Guide*.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Port Security Versus Fabric Binding

Port security and fabric binding are two independent features that can be configured to complement each other. [Table 25-1](#) compares the two features.

Table 25-1 Fabric Binding and Port Security Comparison

Fabric Binding	Port Security
Uses a set of sWWNs and a persistent domain ID.	Uses pWWNs/nWWNs or fWWNs/sWWNs.
Binds the fabric at the switch level.	Binds devices at the interface level.
Authorizes only the configured sWWN stored in the fabric binding database to participate in the fabric.	Allows a preconfigured set of Fibre Channel devices to logically connect to a SAN port. The switch port, identified by a WWN or interface number, connects to a Fibre Channel device (a host or another switch), also identified by a WWN. By binding these two devices, you lock these two ports into a group (or list).
Requires activation on a per VSAN basis.	Requires activation on a per VSAN basis.
Allows specific user-defined switches that are allowed to connect to the fabric, regardless of the physical port to which the peer switch is connected.	Allows specific user-defined physical ports to which another device can connect.
Does not learn about switches that are logging in.	Learns about switches or devices that are logging in if learning mode is enabled.
Cannot be distributed by CFS and must be configured manually on each switch in the fabric.	Can be distributed by CFS.

Port-level checking for xE ports is as follows:

- The switch login uses both port security binding and fabric binding for a given VSAN.
- Binding checks are performed on the port VSAN as follows:
 - E port security binding check on port VSAN
 - TE port security binding check on each allowed VSAN

While port security complements fabric binding, they are independent features and can be enabled or disabled separately.

Fabric Binding Enforcement

To enforce fabric binding, configure the switch world wide name (sWWN) to specify the xE port connection for each switch. Enforcement of fabric binding policies are done on every activation and when the port tries to come up. For a Fibre Channel VSAN, the fabric binding feature requires all sWWNs connected to a switch to be part of the fabric binding active database.

Send comments to nx5000-docfeedback@cisco.com

Configuring Fabric Binding

The fabric binding feature ensures ISLs are only enabled between specified switches in the fabric binding configuration. Fabric binding is configured on a per-VSAN basis.

This section includes the following topics:

- [Configuring Fabric Binding, page 25-3](#)
- [Enabling Fabric Binding, page 25-4](#)
- [About Switch WWN Lists, page 25-4](#)
- [Configuring Switch WWN List, page 25-4](#)
- [About Fabric Binding Activation and Deactivation, page 25-5](#)
- [Activating Fabric Binding, page 25-5](#)
- [Forcing Fabric Binding Activation, page 25-6](#)
- [Copying Fabric Binding Configurations, page 25-6](#)
- [Creating a Fabric Binding Configuration, page 25-7](#)
- [Deleting a Fabric Binding Configuration, page 25-7](#)
- [Copying Fabric Binding to the Configuration File, page 25-8](#)
- [Viewing EFMD Statistics, page 25-8](#)
- [Viewing Fabric Binding Violations, page 25-8](#)
- [Viewing Fabric Binding Active Database, page 25-8](#)
- [Saving Fabric Binding Configurations, page 25-9](#)
- [Clearing the Fabric Binding Statistics, page 25-9](#)
- [Deleting the Fabric Binding Database, page 25-10](#)

Configuring Fabric Binding

To configure fabric binding in each switch in the fabric, perform this task:

-
- | | |
|---------------|---|
| Step 1 | Enable the fabric configuration feature. |
| Step 2 | Configure a list of sWWNs and their corresponding domain IDs for devices that are allowed to access the fabric. |
| Step 3 | Activate the fabric binding database. |
| Step 4 | Copy the fabric binding active database to the fabric binding configuration database. |
| Step 5 | Save the fabric binding configuration. |
| Step 6 | Verify the fabric binding configuration. |
-

Send comments to nx5000-docfeedback@cisco.com

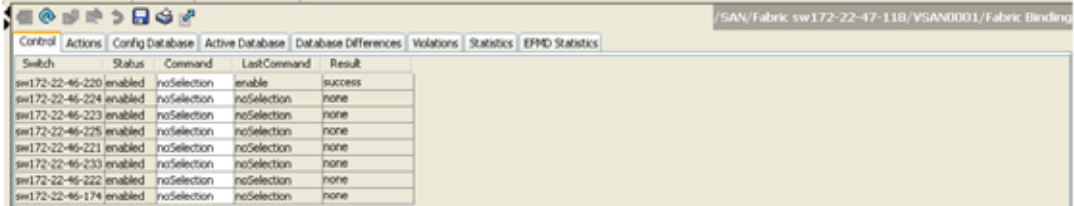
Enabling Fabric Binding

The fabric binding feature must be enabled in each switch in the fabric that participates in the fabric binding. By default, this feature is disabled in Cisco Nexus 5000 Series switches. The configuration and verification commands for the fabric binding feature are only available when fabric binding is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.

To enable fabric binding on any participating switch using Fabric Manager, perform this task:

-
- Step 1** Expand the VSAN with the switches on which you want to enable fabric binding in the Logical Domains pane. Expand **Fabric Binding** (see [Figure 25-1](#)).

Figure 25-1 Fabric Binding Configuration



Switch	Status	Command	LastCommand	Result
sw172-22-46-220	enabled	noSelection	enable	success
sw172-22-46-224	enabled	noSelection	noSelection	none
sw172-22-46-223	enabled	noSelection	noSelection	none
sw172-22-46-225	enabled	noSelection	noSelection	none
sw172-22-46-221	enabled	noSelection	noSelection	none
sw172-22-46-233	enabled	noSelection	noSelection	none
sw172-22-46-222	enabled	noSelection	noSelection	none
sw172-22-46-174	enabled	noSelection	noSelection	none

The **Control** tab is the default tab in the Information pane.

- Step 2** From the Command drop-down list, choose **enable or disable** to enable or disable Fabric Binding on the switch.
- Step 3** Click the **Apply Changes** icon to save your changes.
-

About Switch WWN Lists

A user-specified fabric binding list contains a list of switch WWNs (sWWNs) within a fabric. If an sWWN attempts to join the fabric, and that sWWN is not on 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.

Configuring Switch WWN List

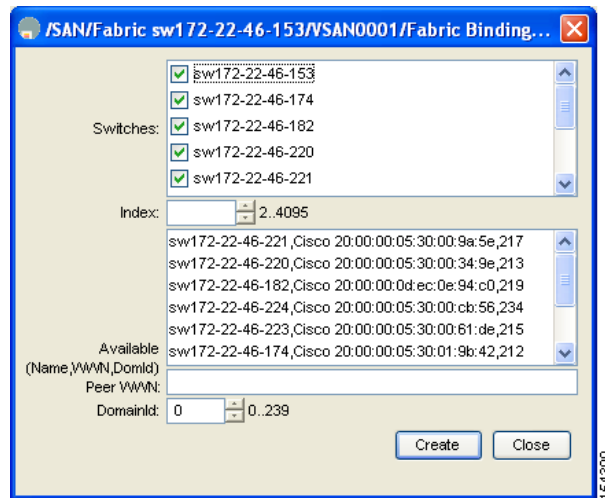
To configure a list of sWWNs and domain IDs for a FICON VSAN using Fabric Manager, perform this task:

-
- Step 1** Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Fabric Binding** (see [Figure 25-1](#)).
- Step 2** Ensure that fabric binding is enabled for the selected VSAN.
- Step 3** Click the **Config Database** tab in the Information pane.
- Step 4** Click **Create Row**.

Send comments to nx5000-docfeedback@cisco.com

You see the Create Config Database dialog box as shown in [Figure 25-2](#).

Figure 25-2 Create Config Database Dialog Box



- Step 5** Select the switches that you want to add.
- Step 6** Add the sWWN and domain ID of a switch to the configured database list.
You can add the sWWN and the domain ID of more than one switches to the configured database list.
- Step 7** Click **Create**.

About Fabric Binding Activation and Deactivation

The fabric binding feature maintains a configuration database (config database) and an active database. The config database is a read-write database that collects the configurations you perform. These configurations are only enforced upon activation. This activation overwrites the active database with the contents of the config database. The active database is read-only and is the database that checks each switch that attempts to log in.

By default, the fabric binding feature is not activated. You cannot activate the fabric binding database on the switch if entries existing in the 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 to forcefully override these situations.



Note

After activation, any already logged in switch that violates the current active database will be logged out, and all switches that were previously denied login because of fabric binding restrictions are reinitialized.

Activating Fabric Binding

To activate, deactivate, or to force fabric bind using Fabric Manager, perform this task:

- Step 1** Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Fabric Binding**.

Send comments to nx5000-docfeedback@cisco.com

Step 2 Click the **Actions** tab in the Information pane (see [Figure 25-3](#)).

Figure 25-3 Fabric Binding Actions Tab

Switch	Action	Enabled	Result	LastChange	CopyActive ToConfig
sw172-22-46-221	activate	false	success	n/a	<input type="checkbox"/>
sw172-22-46-220	activate	false	success	n/a	<input type="checkbox"/>
sw172-22-46-174	forceActivate	false	success	n/a	<input type="checkbox"/>

Step 3 In the Action drop-down list, choose **activate** or **deactivate** or **force activate** for the Fabric Binding on the switch.

Step 4 Click the **Apply Changes** icon to save your changes.
The Enabled column for the switch now displays True.

Forcing Fabric Binding Activation

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed with the activation by using the force option.

To forcefully activate the fabric binding database using Fabric Manager, perform this task:

-
- Step 1** Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Fabric Binding**.
- Step 2** Click the **Actions** tab in the Information pane (see [Figure 25-3](#)).
- Step 3** In the Action drop-down list, choose **forceActivate** for the VSAN(s) for which you want to activate fabric binding.
- Step 4** Click **Apply Changes** to activate fabric binding.
The Enabled column for the switch now displays True.
-

Copying Fabric Binding Configurations

When you copy the fabric binding configuration, the config database is saved to the running configuration.

To copy the active database to the config database using Fabric Manager, perform this task:

-
- Step 1** Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Fabric Binding**.
- Step 2** Click the **Actions** tab in the Information pane.
- Step 3** Check the **Copy Active to Config** check box.

Send comments to nx5000-docfeedback@cisco.com

Step 4 Click the **Apply Changes** icon to save your changes.

Creating a Fabric Binding Configuration

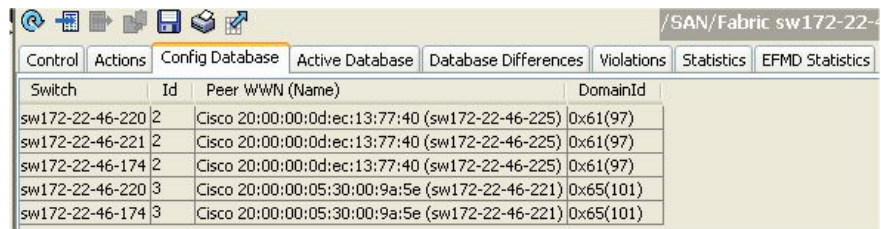
To create a fabric binding configuration using Fabric Manager, perform this task:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Fabric Binding**.

Step 2 Click the **Config Database** tab in the Information pane.

You see the information as shown in [Figure 25-4](#).

Figure 25-4 Fabric Binding Database Configuration



Switch	Id	Peer WWN (Name)	DomainId
sw172-22-46-220	2	Cisco 20:00:00:0d:ec:13:77:40 (sw172-22-46-225)	0x61(97)
sw172-22-46-221	2	Cisco 20:00:00:0d:ec:13:77:40 (sw172-22-46-225)	0x61(97)
sw172-22-46-174	2	Cisco 20:00:00:0d:ec:13:77:40 (sw172-22-46-225)	0x61(97)
sw172-22-46-220	3	Cisco 20:00:00:05:30:00:9a:5e (sw172-22-46-221)	0x65(101)
sw172-22-46-174	3	Cisco 20:00:00:05:30:00:9a:5e (sw172-22-46-221)	0x65(101)

Step 3 Click **Insert Row**.

You see the Create Config Database dialog box (see [Figure 25-2](#)).

Step 4 Select switches, choose an index, and indicate the peer WWN and the Domain ID.

Step 5 Click **Create** to create the fabric binding database configuration.

When you save the fabric binding configuration, the config database and the active database are both saved to the startup configuration and are available after a reboot.

Deleting a Fabric Binding Configuration

To delete a fabric binding configuration using Fabric Manager, perform this task:

Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Fabric Binding**.

Step 2 Click the **Config Database** tab in the Information pane.

You see the information shown in [Figure 25-4](#).

Step 3 Click in the row for the VSAN for which you want to delete the fabric binding configuration.

Step 4 Click **Delete Row** to delete the fabric binding configuration.

Send comments to nx5000-docfeedback@cisco.com

Copying Fabric Binding to the Configuration File

To copy the active fabric binding to the configuration file using Fabric Manager, perform this task:

-
- Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Fabric Binding**.
 - Step 2 Click the **Actions** tab in the Information pane (see [Figure 25-3](#)).
 - Step 3 Check the **CopyActive ToConfig** check box for the VSAN(s) for which you want to copy fabric binding.
 - Step 4 Click the **Apply Changes** icon to copy the fabric binding.
-



Caution

You cannot deactivate or disable fabric binding in a FICON-enabled VSAN.

Viewing EFMD Statistics

To view EFMD statistics using Fabric Manager, perform this task:

-
- Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Fabric Binding**.
 - Step 2 Click the **EFMD Statistics** tab.
- You see the statistics information.
-

Viewing Fabric Binding Violations

To view fabric binding violations using Fabric Manager, perform this task:

-
- Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Fabric Binding**.
 - Step 2 Click the **Violations** tab.
- You see the violations information.
-


Viewing Fabric Binding Active Database

To view the fabric binding active database using Fabric Manager, perform this task:

-
- Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Fabric Binding**.
 - Step 2 Click the **Active Database** tab.
- You see the active database information as shown in [Figure 25-5](#).
-

Send comments to nx5000-docfeedback@cisco.com

Figure 25-5 Fabric Binding Active Database



Saving Fabric Binding Configurations

When you save the fabric binding configuration, the config database and the active database are both saved to the startup configuration and are available after a reboot.

To save the fabric binding configuration using Fabric Manager, perform this task:

-
- Step 1** Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Fabric Binding**.
 - Step 2** Click the **Actions** tab (see [Figure 25-3](#)).
 - Step 3** Check the **Copy Active to Config** check box to copy the active database to the config database.
If the configured database is empty, this action is not successful.
 - Step 4** Click the **Database Differences** tab to compare the database with the Config or Active database to view the differences between the active database and the config database.
-

Clearing the Fabric Binding Statistics

To clear all existing statistics from the fabric binding database for a specified VSAN using Fabric Manager, perform this task:

-
- Step 1** Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Fabric Binding**.
 - Step 2** Click the **Statistics** tab in the Information pane.
You see the statistics in the Information pane.
 - Step 3** Check the **Clear** check box.
 - Step 4** Click the **Apply Changes** icon to save your changes.
-

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Deleting the Fabric Binding Database

To delete the configured database for a specified VSAN using Fabric Manager, perform this task:

-
- Step 1 Expand a VSAN with fabric binding in the Logical Domains pane. Expand **Fabric Binding**.
 - Step 2 Click the **Config Database** tab in the Information pane.
 - Step 3 Select the database that you want to delete.
 - Step 4 Click **Delete Row**.
-

Default Settings

Table 25-2 lists the default settings for the fabric binding feature.

Table 25-2 Default Fabric Binding Settings

Parameters	Default
Fabric binding	Disabled



Configuring Fabric Configuration Servers

This chapter describes the Fabric Configuration Server (FCS) feature provided in the Cisco Nexus 5000 Series switches. It includes the following sections:

- [Information About FCS, page 26-1](#)
- [Displaying FCS Discovery, page 26-3](#)
- [Displaying FCS Elements, page 26-3](#)
- [Creating an FCS Platform, page 26-4](#)
- [Displaying FCS Fabric Ports, page 26-5](#)
- [Default Settings, page 26-6](#)

Information 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 and F ports) and their attached N ports.
- Platform object—A set of nodes may be defined as a platform object to make it a single manageable entity. These nodes are end-devices (host systems, storage subsystems) attached to the fabric. Platform objects reside at the edge switches of the fabric.

Each object has its own set of attributes and values. A null value may also be defined for some attributes.

In the Cisco Nexus 5000 Series switch environment, a fabric may consist of multiple VSANs. One instance of the FCS is present per VSAN.

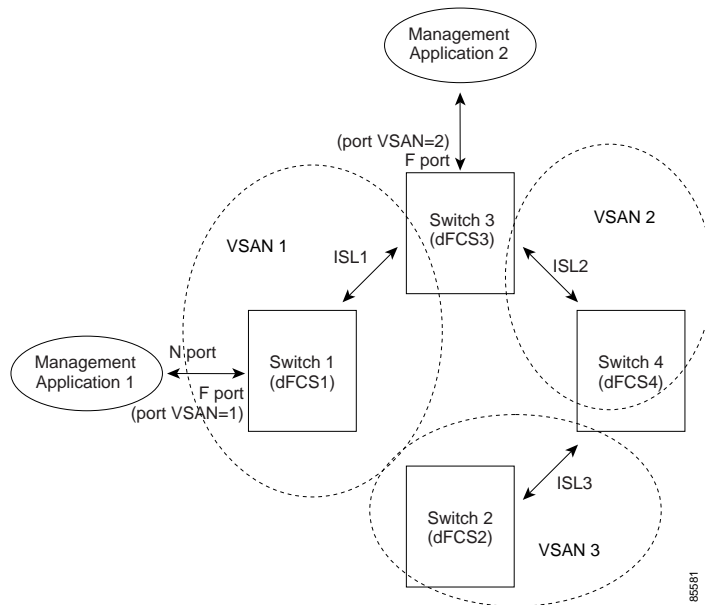
FCS supports the discovery of virtual devices. The **fcs virtual-device-add** command, entered in FCS configuration submode, allows you to discover virtual devices in a particular VSAN or in all VSANs.

If you have attached a management application to a switch, all the frames directed towards the FCS in the switch are part of the port VSAN in the switch port (F port). Your view of the management application is limited only to this VSAN. However, information about other VSANs that this switch is part of can be obtained either through the SNMP or CLI.

Send comments to nx5000-docfeedback@cisco.com

In [Figure 26-1](#) Management Application 1 (M1) is connected through an F port with port VSAN ID 1, and Management Application 2 (M2) is connected through an F port with port VSAN ID 2. M1 can query the FCS information of switches S1 and S3, and M2 can query switches S3 and S4. Switch S2 information is not known to both of them. FCS operations can be done only on those switches that are visible in the VSAN. M2 can send FCS requests only for VSAN 2 even though S3 is also a part of VSAN 1.

Figure 26-1 FCSs in a VSAN Environment



FCS Characteristics

FCSs have the following characteristics:

- Support network management including the following:
 - N port management application can query and obtain information about fabric elements.
 - SNMP manager can use the FCS management information base (MIB) to start discovery and obtain information about the fabric topology.
- Support TE ports in addition to the standard F and E ports.
- Can maintain a group of nodes with a logical name and management address when a platform registers with it. FCSs maintain a backup of all registrations in secondary storage and update it with every change. When a restart or switchover happens, FCSs retrieve the secondary storage information and rebuild its database.
- SNMP manager can query FCSs for all IEs, ports, and platforms in the fabric.

Send comments to nx5000-docfeedback@cisco.com

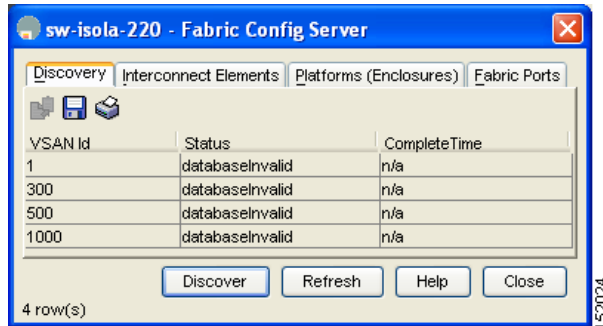
Displaying FCS Discovery

To display FCS discovery information using Device Manager, perform this task:

Step 1 Choose **FC > Advanced > Fabric Config Server**.

You see the Fabric Config Server dialog box as shown in [Figure 26-2](#).

Figure 26-2 Fabric Config Server Dialog Box



Step 2 Click the **Discovery** tab.

Step 3 Click **Discover** to rediscover the fabric, or click **Refresh** to update the display.

Displaying FCS Elements

To display FCS interconnect element information using Device Manager, perform this task:

Step 1 Choose **FC > Advanced > Fabric Config Server**.

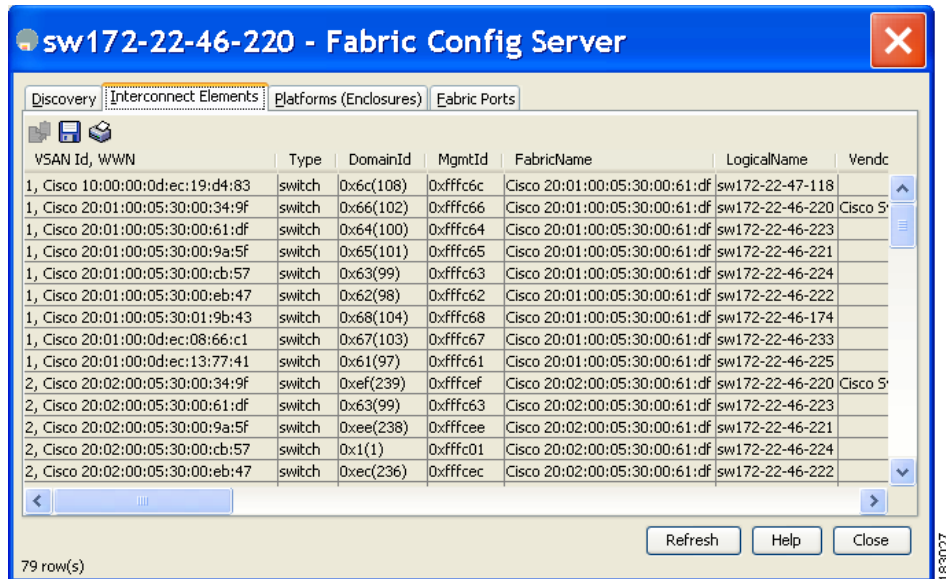
You see the Fabric Config Server dialog box.

Step 2 Click the **Interconnect Elements** tab.

You see the dialog box shown in [Figure 26-3](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 26-3 FCS Interconnect Elements Tab



Step 3 Click **Close** to close the dialog box.

Creating an FCS Platform

To create an FCS platform using Device Manager, perform this task:

Step 1 Choose **FC > Advanced > Fabric Config Server**.

You see the Fabric Config Server dialog box.

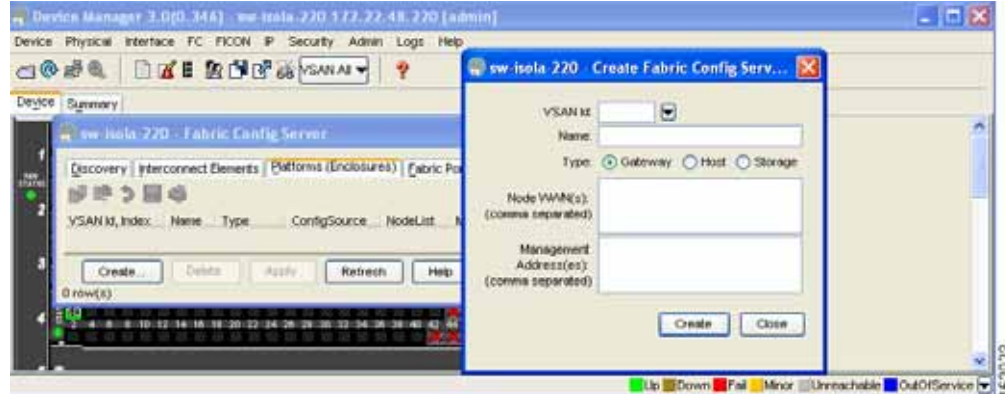
Step 2 Click the **Platforms (Enclosures)** tab.

Step 3 Click **Create**.

You see the Create Fabric Config Server dialog box as shown in [Figure 26-4](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 26-4 Create Fabric Config Server Dialog Box



- Step 4** Enter the VSAN ID, or choose the ID from the drop-down list of available VSAN IDs.
- Step 5** Enter the Fabric Configuration Server name in the Name field.
- Step 6** Choose the type of server (**Gateway**, **Host**, **Storage**).
- Step 7** Enter the WWNs for the server.
- Step 8** Enter the management addresses for the server.
- Step 9** Click **Create** to create the server.

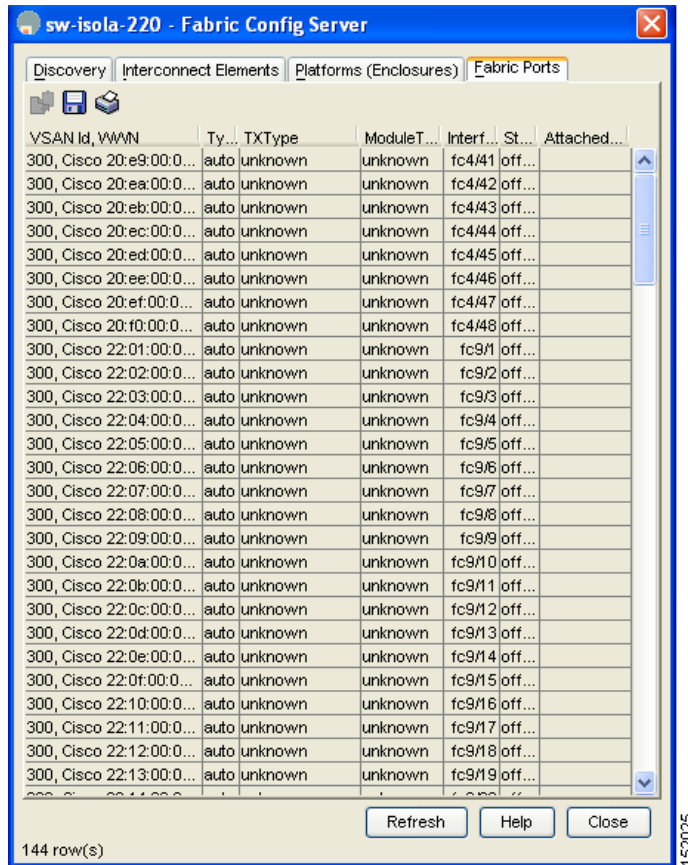
Displaying FCS Fabric Ports

To display FCS discovery information using Device Manager, perform this task:

- Step 1** Choose **FC > Advanced > Fabric Config Server**.
You see the Fabric Config Server dialog box.
- Step 2** Click the **Fabric Ports** tab.
You see a list of fabric ports as shown in [Figure 26-5](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 26-5 FCS Fabric Ports Tab



Step 3 Click **Refresh** to update the display.

Default Settings

Table 26-1 lists the default FCS settings.

Table 26-1 Default FCS Settings

Parameters	Default
Global checking of the platform name	Disabled
Platform node type	Unknown



Configuring Port Tracking

Cisco Nexus 5000 Series switches offer the port tracking feature on physical Fibre Channel interfaces (but not on virtual Fibre Channel interfaces). This feature uses information about the operational state of the link to initiate a failure in the link that connects the edge device. This process of converting the indirect failure to a direct failure triggers a faster recovery process towards redundant links. When enabled, the port tracking feature brings down the configured links based on the failed link and forces the traffic to be redirected to another redundant link.

This chapter includes the following sections:

- [Information About Port Tracking, page 27-1](#)
- [Configuring Port Tracking, page 27-2](#)
- [Default Port Tracking Settings, page 27-7](#)


Information About Port Tracking

Generally, hosts can instantly recover from a link failure on a link that is immediately (direct link) connected to a switch. However, recovering from an indirect link failure between switches in a WAN or MAN fabric with a keepalive mechanism is dependent on several factors such as the timeout values (TOVs) and on registered state change notification (RSCN) information (see the [“Fibre Channel Timeout Values”](#) section on page 22-1 and [“About RSCN Information”](#) section on page 19-5).

In [Figure 27-1](#), when the direct link 1 to the host fails, recovery can be immediate. However, when the ISL 2 fails between the two switches, recovery depends on TOVs, RSCNs, and other factors.

Send comments to nx5000-docfeedback@cisco.com

Figure 27-1 Traffic Recovery Using Port Tracking



The port tracking feature monitors and detects failures that cause topology changes and brings down the links connecting the attached devices. When you enable this feature and explicitly configure the linked and tracked ports, the switch software monitors the tracked ports and alters the operational state of the linked ports on detecting a link state change.

The following terms are used in this chapter:

- **Tracked ports**—A port whose operational state is continuously monitored. The operational state of the tracked port is used to alter the operational state of one or more ports. Fibre Channel, VSAN, SAN port channel, or a Gigabit Ethernet port can be tracked. Generally, ports in E and TE port modes can also be F ports.
- **Linked ports**—A port whose operational state is altered based on the operational state of the tracked ports. Only physical Fibre Channel ports can be linked ports.

Port tracking has the following features:

- The application brings the linked port down when the tracked port goes down. When the tracked port recovers from the failure and comes back up again, the linked port is also brought up automatically (unless otherwise configured).
- You can forcefully continue to keep the linked port down, even though the tracked port comes back up. In this case, you must explicitly bring up the linked port when required.

Configuring Port Tracking

Before configuring port tracking, consider the following guidelines:

- Verify that the tracked ports and the linked ports are on the same Cisco switch.
- Be aware that the linked port is automatically brought down when the tracked port goes down.
- Do not track a linked port back to itself (for example, Port fc2/2 to Port fc2/4 and back to Port fc2/2) to avoid recursive dependency.

This section includes the following topics:

- [Enabling Port Tracking, page 27-3](#)
- [About Configuring Linked Ports, page 27-3](#)
- [Operationally Binding a Tracked Port, page 27-4](#)
- [About Tracking Multiple Ports, page 27-5](#)

Send comments to nx5000-docfeedback@cisco.com

- [Tracking Multiple Ports, page 27-6](#)
- [About Monitoring Ports in a VSAN, page 27-6](#)
- [Monitoring Ports in a VSAN, page 27-6](#)
- [About Forceful Shutdown, page 27-6](#)
- [Forcefully Shutting Down a Tracked Port, page 27-6](#)

Enabling Port Tracking

The port tracking feature is disabled by default in Cisco Nexus 5000 Series switches. When you enable this feature, port tracking is globally enabled for the entire switch.

To configure port tracking, enable the port tracking feature and configure the linked ports for the tracked port.

To enable port tracking with Fabric Manager, perform this task:

- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **Port Tracking**. The port tracking information is displayed in the Information pane as shown in [Figure 27-2](#). The default tab is the Controls tab.

Figure 27-2 Port Tracking

Switch	Status	Command	LastCommand	Result
sw172-22-46-220	enabled	noSelection	noSelection	none
sw172-22-46-224	disabled	noSelection	noSelection	none
sw172-22-46-223	enabled	noSelection	noSelection	none
sw172-22-46-221	disabled	noSelection	noSelection	none
sw172-22-46-222	enabled	noSelection	noSelection	none
sw172-22-46-223	disabled	noSelection	noSelection	none
sw172-22-46-225	disabled	noSelection	noSelection	none
sw172-22-46-174	enabled	noSelection	noSelection	none

- Step 2** In the Command column, choose **enable** or **disable** for the port tracking. Depending on your selection the corresponding entry in the Status column changes.
- Step 3** Click the **Apply Changes** icon to save your changes. The entry in the Result column changes to success.

About Configuring Linked Ports

You can link ports using one of two methods:

- Operationally binding the linked ports to the tracked port (default).
- Continuing to keep the linked port down forcefully, even if the tracked port has recovered from the link failure.

Send comments to nx5000-docfeedback@cisco.com

Operationally Binding a Tracked Port

When you configure the first tracked port, operational binding is automatically in effect. When you use this method, you have the option to monitor multiple ports or monitor ports in one VSAN.

To operationally bind a tracked port, perform this task:

- Step 1** In the Physical Attributes pane, expand **Switches > Interfaces**, and then choose **Port Tracking**. The port tracking information is displayed in the Information pane. The default tab is the Controls tab.

Figure 27-3 Port Tracking Controls Tab



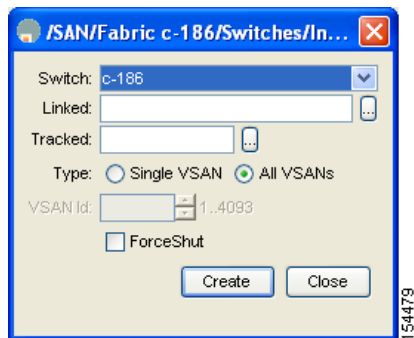
Switch	Status	Command	LastCommand	Result
sw172-22-46-220	enabled	noSelection	noSelection	none
sw172-22-46-224	disabled	noSelection	noSelection	none
sw172-22-46-223	enabled	noSelection	noSelection	none
sw172-22-46-221	disabled	noSelection	noSelection	none
sw172-22-46-222	enabled	noSelection	noSelection	none
sw172-22-46-233	disabled	noSelection	noSelection	none
sw172-22-46-225	disabled	noSelection	noSelection	none
sw172-22-46-174	enabled	noSelection	noSelection	none

- Step 2** Click the **Dependencies** tab.

- Step 3** Click **Create Row**.

You see the Create Port Tracking Dependencies dialog box as shown in [Figure 27-4](#).

Figure 27-4 Create Port Tracking Dependencies Dialog Box

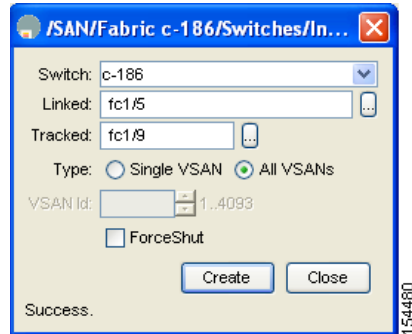


- Step 4** Select the switch whose ports you want to track by choosing a switch from the drop-down list.
- Step 5** Select the linked ports that should be bound to the tracked ports by clicking the browse button and choosing from the list.
- Step 6** Click the **Single VSAN** radio button if you want to track these ports only in one VSAN or click the **All VSANs** radio button if you want to track these ports in all the available VSANs. See [“About Monitoring Ports in a VSAN”](#) section on page 27-6 for details.
- Step 7** If you chose Single VSAN in the previous step, enter the ID of the VSAN where these ports will be monitored.
- Step 8** Check the **Forceshut** check box if you want to forcefully shut down the tracked port. See [“About Forceful Shutdown”](#) section on page 27-6 for details.
- Step 9** Click **Create** to proceed with creating this dependency.

Send comments to nx5000-docfeedback@cisco.com

If tracking is established, you see **Success** in the lower left corner of the dialog box as shown in [Figure 27-5](#).

Figure 27-5 Successful Port Tracking Established



Step 10 Click **Close** to close the dialog box.

The following example shows how to enable port tracking for specific interfaces:


```
switch# configure
switch(config)# interface fc 2/4
switch(config-if)# port-track interface san-port-channel 2
```

About Tracking Multiple Ports

You can control the operational state of the linked port based on the operational states of multiple tracked ports. When more than one tracked port is associated with a linked port, the operational state of the linked port will be set to down only if all the associated tracked ports are down. Even if one tracked port is up, the linked port will stay up.

In [Figure 27-6](#), only if both ISLs 2 and 3 fail, will the direct link 1 be brought down. Direct link 1 will not be brought down if either 2 or 3 are still functioning as desired.

Figure 27-6 Traffic Recovery Using Port Tracking



Send comments to nx5000-docfeedback@cisco.com

Tracking Multiple Ports

To track multiple ports, see the “[Operationally Binding a Tracked Port](#)” section on page 27-4.

The following example shows how to enable port tracking for multiple interfaces:

```
switch# configure
switch(config)# interface fc 2/3
switch(config-if)# port-track interface fc 2/3
switch(config-if)# port-track interface san-port-channel 2
```

About Monitoring Ports in a VSAN

You can optionally configure one VSAN from the set of all operational VSANs on the tracked port with the linked port by specifying the required VSAN. This level of flexibility provides higher granularity in tracked ports. In some cases, when a tracked port is a TE port, the set of operational VSANs on the port can change dynamically without bringing down the operational state of the port. In such cases, the port VSAN of the linked port can be monitored on the set of operational VSANs on the tracked port.

If you configure this feature, the linked port is up only when the VSAN is up on the tracked port.



Tip

The specified VSAN does not have to be the same as the port VSAN of the linked port.

Monitoring Ports in a VSAN

To monitor a tracked port in a specific VSAN, see “[Operationally Binding a Tracked Port](#)” section on page 27-4.

About Forceful Shutdown

If a tracked port flaps frequently, then tracking ports using the operational binding feature may cause frequent topology change. In this case, you may choose to keep the port in the down state until you are able to resolve the reason for these frequent flaps. Keeping the flapping port in the down state forces the traffic to flow through the redundant path until the primary tracked port problems are resolved. When the problems are resolved and the tracked port is back up, you can explicitly enable the interface.



Tip

If you configure this feature, the linked port continues to remain in the shutdown state even after the tracked port comes back up. You must explicitly remove the forced shut state (by administratively bringing up this interface) of the linked port once the tracked port is up and stable.

Forcefully Shutting Down a Tracked Port

To forcefully shut down a tracked port, see “[Operationally Binding a Tracked Port](#)” section on page 27-4.

Send comments to nx5000-docfeedback@cisco.com

Default Port Tracking Settings

Table 27-1 lists the default settings for port tracking parameters.

Table 27-1 *Default Port Tracking Parameters*

Parameters	Default
Port tracking	Disabled
Operational binding	Enabled along with port tracking

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)



Network Monitoring

The primary purpose of Fabric Manager is to manage the network. In particular, SAN discovery and network monitoring are two of its key network management capabilities.

This chapter contains the following sections:

- [Information About SAN Discovery and Topology Mapping, page 28-1](#)
- [Health and Event Monitoring, page 28-4](#)

Information About SAN Discovery and Topology Mapping

Fabric Manager provides extensive SAN discovery, topology mapping, and information viewing capabilities. Fabric Manager collects information on the fabric topology through SNMP queries to the switches connected to it. Fabric Manager recreates a fabric topology, presents it in a customizable map, and provides inventory and configuration information in multiple viewing options.

This section includes the following topics:

- [Device Discovery, page 28-1](#)
- [Topology Mapping, page 28-2](#)
- [Inventory Management, page 28-3](#)
- [Viewing Logs from Device Manager, page 28-4](#)

Device Discovery

Once Fabric Manager is invoked, a SAN discovery process begins. Using information polled from a seed switch, including Name Server registrations, Fibre Channel Generic Services (FC-GS), Fabric Shortest Path First (FSPF), and SCSI-3, Fabric Manager automatically discovers all devices and interconnects on one or more fabrics. All available switches, host bus adapters (HBAs), and storage devices are discovered. Cisco SAN switches use Fabric-Device Management Interface (FMDI) to retrieve HBA model, serial number and firmware version, and host operating-system type and version discovery without host agents. Fabric Manager obtains this information through SNMP queries to each switch. The device information discovered includes device names, software revision levels, vendor, ISLs, SAN port channels, and VSANs.

Send comments to nx5000-docfeedback@cisco.com

Topology Mapping

Fabric Manager is built upon a topology representation of the fabric. Fabric Manager provides an accurate view of multiple fabrics in a single window by displaying topology maps based on device discovery information. You can modify the topology map icon layout with an easy-to-use, drag-and-drop interface. The topology map visualizes device interconnections, highlights configuration information such as zones, VSANs, and ISLs exceeding utilization thresholds. The topology map also provides a visual context for launching command-line interface (CLI) sessions, configuring SAN port channels, and opening device managers.

Using the Topology Map

The Fabric Manager topology map can be customized to provide a view into the fabric that varies from showing all switches, end devices, and links, to showing only the core switches with single bold lines for any multiple links between switches. Use the icons along the left side of the topology map to control these views or right-click anywhere in the topology map to access the map controls.

You can zoom in or out on the topology map to see an overview of the SAN or focus on an area of importance. You can also open an overview window that shows the entire fabric. From this window, you can right-click and draw a box around the area you want to view in the main topology map view.

Another way to limit the scope of the topology display is to select a fabric or VSAN from the Logical Domains pane. The topology map displays only that fabric or VSAN.

Moving the mouse pointer over a link or switch provides a simple summary of that SAN component, along with a status indication. Right-clicking on the component brings up a pop-up menu. You can view the component in detail or access configuration or test features for that component.

Double-click a link to bring link status and configuration information to the information pane.
Double-click a switch to bring up Device Manager for that switch.

Saving a Customized Topology Map Layout

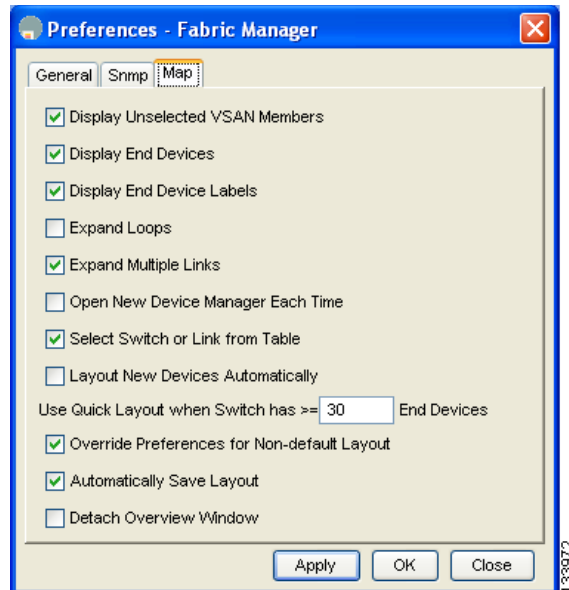
Changes made to the topology map can be saved so that the customized view is available any time you open the Fabric Manager Client for that fabric.

To save the customized layout using Fabric Manager, perform this task:

-
- Step 1** Choose **File > Preferences** to open the Fabric Manager preferences dialog box.
 - Step 2** Click the **Map** tab and check the **Automatically Save Layout** check box to save any changes to the topology map (See [Figure 28-1](#)).

Send comments to nx5000-docfeedback@cisco.com

Figure 28-1 Fabric Manager Preferences



Step 3 Click **Apply**, then **OK** to save this change.

Using Enclosures with Fabric Manager Topology Maps

Because not all devices are capable of responding to FC-GS-3 requests, different ports of a single server or storage subsystem may be displayed as individual end devices on the topology map. See the [“Modifying the Device Grouping” section on page 5-20](#) to group these ports into a single enclosure for Fabric Manager.

Choosing **Alias->Enclosure** displays hosts and storage elements in the Information pane. This is a shortcut to naming enclosures. To use this shortcut, highlight each row in the host or storage table that you want grouped in an enclosure, and then choose **Alias -> Enclosure**. This automatically sets the enclosure names of each selected row with the first token of the alias.

Mapping Multiple Fabrics

To log into multiple fabrics, the same username and password must be used. The information for both fabrics is displayed, with no need to select a seed switch. To see details of a fabric, click the tab for that fabric at the bottom of the Fabric pane, or double-click the fabric’s cloud icon.

To continuously manage a fabric using Fabric Manager, follow the instructions in the [“Managing a Fabric Manager Server Fabric” section on page 3-3](#).

Inventory Management

The Information pane in Fabric Manager shows inventory, configuration, and status information for all switches, links, and hosts in the fabric. Inventory management includes vendor name and model, and software or firmware versions. Select a fabric or VSAN from the Logical Domains pane, and then choose

Send comments to nx5000-docfeedback@cisco.com

the **Summary** tab in the Information pane to get a count of the number of VSANS, switches, hosts, and storage elements in the fabric. See the “[Fabric Manager Client Quick Tour](#)” section on page 5-6 for more information on the Fabric Manager user interface.

Using the Inventory Tab from Fabric Manager Web Server

If you have configured Fabric Manager Web Server, you can launch this application and access the Inventory tab to see a summary of the fabrics managed by the Fabric Manager Server. The Inventory tab shows an inventory of the selected SAN, fabric, or switch.

See the *Cisco MDS 9000 Family Fabric Manager Configuration Guide* for additional information on how to configure and use Fabric Manager Web Server.

To view system messages remotely using Fabric Manager Web Server, perform this task:

-
- Step 1 Point your browser at the Fabric Manager Web Server.
 - Step 2 Click the **Events** tab, and then choose **Details** to view the system messages.

The columns in the events table are sortable. In addition, you can use the Filter button to limit the scope of messages within the table.

Viewing Logs from Device Manager

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

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

Health and Event Monitoring

Fabric Manager works with the Cisco SAN switches to show the health and status of the fabric and switches. Information about the fabric and its components is gathered from multiple sources, including Online System Health Management, Call Home, system messages, and SNMP notifications. This information is then made available from multiple menus on Fabric Manager or Device Manager.

Fabric Manager Events Tab

The Fabric Manager Events tab, available from the topology window, displays the events Fabric Manager received from sources within the fabric. These sources include SNMP events, RMON events, system messages, and system health messages. The Events tab shows a table of events, including the event name, the source and time of the event, a severity level, and a description of the event. The table is sortable by any of these column headings.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Event Information in Fabric Manager Web Server Reports

The Fabric Manager web server client displays collections of information gathered by the Performance Manager. This information includes events sent to the Fabric Manager Server from the fabric. To open these reports, choose **Performance Manager > Reports**. This opens the web client in a web browser and displays a summary of all fabrics monitored by the Fabric Manager Server. Choose a fabric, and then choose the **Events** tab to see a summary or detailed report of the events that have occurred in the selected fabric. The summary view shows how many switches, ISLs, hosts, or storage elements are down on the fabric and how many warnings have been logged for that SAN entity. The detailed view shows a list of all events that have been logged from the fabric and can be filtered by severity, time period, or type.

Events in Device Manager

Device Manager displays the events when you choose **Logs > Events**. Device Manager can display the current list of events or an older list of events that has been stored on the Fabric Manager host. The event table shows details on each event, including time, source, severity, and a brief description of the event.

Send comments to nx5000-docfeedback@cisco.com



Performance Manager

The primary purpose of Fabric Manager is to manage the network. A key management capability is network performance monitoring. This chapter includes the following topics:

- [Information About Performance Manager, page 29-1](#)
- [Flow Statistics Configuration, page 29-4](#)

Information About Performance Manager

Performance Manager gathers network device statistics historically and provides this information graphically using a web browser. It presents recent statistics in detail and older statistics in summary. Performance Manager also integrates with external tools such as Cisco Traffic Analyzer.

The Performance Manager has three operational stages:

- **Definition**—The Flow Wizard sets up flows in the switches.
- **Collection**—The Web Server Performance Collection screen collects information on desired fabrics.
- **Presentation**—Generates web pages to present the collected data through Fabric Manager Web Server.

Performance Manager can collect statistics for ISLs, hosts, storage elements, and configured flows. Flows are defined based on a host-to-storage (or storage-to-host) link. Performance Manager gathers statistics from across the fabric based on collection configuration files. These files determine which SAN elements and SAN links Performance Manager gathers statistics for. Based on this configuration, Performance Manager communicates with the appropriate devices (switches, hosts, or storage elements) and collects the appropriate information at fixed five-minute intervals.

Performance Manager uses a round-robin database to hold the statistical data collected from the fabric. This data is stored based on the configured parameters in the collection configuration file. At each polling interval, Performance Manager gathers the relevant statistics and stores them in the round-robin database. This database is a fixed size and will not grow beyond its preset limits.

Performance Manager creates a series of archived data to hold summarized information present in the real-time round-robin database. This archived data is used to generate daily, weekly, monthly, and yearly consolidated reports. In this way, Performance Manager maintains significant historical data without the cost of an ever-increasing database size.



Note

You must restart Performance Manager if you change the user credentials on Fabric Manager Server.

Send comments to nx5000-docfeedback@cisco.com

This section includes the following topics:

- [Data Interpolation, page 29-2](#)
- [Data Collection, page 29-2](#)
- [Using Performance Thresholds, page 29-2](#)
- [Flow Setup Wizards, page 29-3](#)

Data Interpolation

One of the unique features of Performance Manager is its ability to interpolate data when statistical polling results are missing or delayed. Other performance tools may store the missing data point as zero, but this can distort historical trending. Performance Manager interpolates the missing data point by comparing the data point that preceded the missing data and the data point stored in the polling interval after the missing data. This maintains the continuity of the performance information.

Data Collection

One year's collection of data for two variables (Rx and Tx bytes) requires a round-robin database (rrd) file size of 76 KB. If errors and discards are also collected, the rrd file size becomes 110K. The following are the default internal values:

- 600 samples of 5 minutes (2 days and 2 hours)
- 700 samples of 30 minutes (12.5 days)
- 775 samples of 2 hours (50 days)
- 300 samples of 1 day

A 1000-port SAN requires 110 MB for a year's collection of historical data that includes errors and discards. If there were 20 switches in this SAN with equal distribution of fabric ports, about two to three SNMP packets per switch would be sent every 5 minutes for a total of about 100 request or response SNMP packets required to monitor the data.

Flows are more difficult to predict storage space requirements for, because of their variable counter requests. Each extra flow adds an additional 76 KB.

**Note**

Performance Manager does not collect statistics on nonmanageable and non-Cisco switches. Loop devices (FL/NL) are not collected.

Using Performance Thresholds

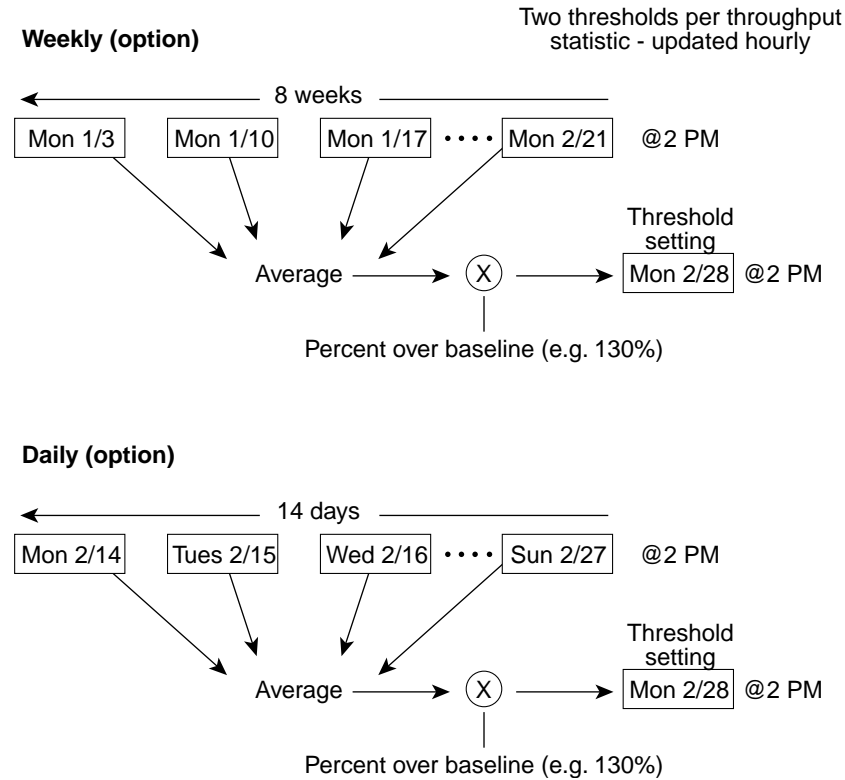
The Performance Manager Configuration Wizard allows you to set up two thresholds that will trigger events when the monitored traffic exceeds the percent utilization configured. These event triggers can be set as either Critical or Warning events that are reported on the Fabric Manager web client Events browser page.

Absolute value thresholds apply directly to the statistics gathered. These statistics, as a percent of the total link capacity, are compared to the percent utilization configured for the threshold type. If the statistics exceed either configured threshold, an event is shown on the Fabric Manager web client Events tab.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Baseline thresholds create a threshold that adapts to the typical traffic pattern for each link for the same time window each day, week, or every two weeks. Baseline thresholds are set as a percent of the average (110 percent to 500 percent), where 100 percent equals the calculated weighted average. [Figure 29-1](#) shows an example of setting a baseline threshold for a weekly or daily option.

Figure 29-1 Baseline Threshold Example



The threshold is set for Monday at 2 p.m. The baseline threshold is set at 130 percent of the average for that statistic. The average is calculated from the statistics value that occurred at 2 p.m. on Monday, for every prior Monday (for the weekly option) or the statistics value that occurred at 2 p.m. on each day, for every prior day (for the daily option).

Flow Setup Wizards

The Performance Manager Flow and Performance Manager Setup wizards greatly simplify configuration. You only need to select the categories of statistics to capture and the wizards provide a list of flows and links to monitor. You can remove entries if desired, or just accept the provided list and start data collection. Statistics for host and storage links are not associated with a specific port on a switch, so you do not lose long term statistics if a connection is moved to a different port.

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Flow Statistics Configuration

Flow statistics count the ingress traffic in the aggregated statistics table. You can collect two kinds of statistics:

- Aggregated flow statistics to count the traffic for a VSAN.
- Flow statistics to count the traffic for a source and destination ID pair in a VSAN.

This section includes the following topics:

- [About Flow Statistics, page 29-4](#)
- [Counting Flow Statistics, page 29-4](#)

About Flow Statistics

If you enable flow counters, you can enable a maximum of 1024 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.

[Table 29-1](#) explains the Flow Type radio button that defines the type of traffic monitored.

Table 29-1 Performance Manager Flow Types

Flow type	Description
Host->Storage	Unidirectional flow, which monitors data from the host to the storage element
Storage->Host	Unidirectional flow, which monitors data from the storage element to the host
Both	Bidirectional flow, which monitors data to and from the host and storage elements.

Counting Flow Statistics

To count the flow statistics for an active zone set using Fabric Manager, perform this task:

-
- Step 1** Expand **End Devices**, and then choose **Flow Statistics** in the Physical Attributes pane. You see the Flow Statistics as shown in [Figure 29-2](#).

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Figure 29-2 Flow Statistics in Fabric Manager

Switch	Module, Index	Type	Vsanid	Destid	Srcid	Mask	Frames	Bytes	CreationTime
sw172-22-46-225	1, 1	vsanid destid srcid	4001	df0007	df001f	fffff	0	0	2007/03/29-14:00:04
sw172-22-46-223	1, 2	vsanid destid srcid	4001	ef0003	ec0007	fffff	0	0	2007/03/29-14:04:06
sw172-22-46-233	1, 1	vsanid destid srcid	1ee0006	ed0004	fffff	fffff	8.523K	505.100K	2007/04/02-11:42:57
sw172-22-46-221	4, 1	vsanid destid srcid	1ee0006	6401e4	fffff	fffff	631	37.440K	2007/04/04-12:44:05
sw172-22-46-220	1, 72	vsanid destid srcid	4001	ea019b	ef0003	fffff	0	0	2007/03/31-19:50:14
sw172-22-46-224	1, 1	vsanid destid srcid	4001	ef0003	ea019f	fffff	712	104.612K	2007/02/21-08:03:55
sw172-22-46-174	10, 1	vsanid destid srcid	1ee0006	ee00ef	fffff	fffff	755	47.880K	2007/04/04-12:44:10
sw172-22-46-233	1, 2	vsanid destid srcid	1ef000f	ed0004	fffff	fffff	1.031K	59.050K	2007/04/02-11:42:57
sw172-22-46-223	1, 3	vsanid destid srcid	1ee0006	650101	fffff	fffff	16.231K	682.832K	2007/04/02-11:43:05
sw172-22-46-221	4, 2	vsanid destid srcid	1ef000f	6401e4	fffff	fffff	998.944M	2067.213G	2007/04/04-12:44:05
sw172-22-46-220	1, 77	vsanid destid srcid	4001	ea019f	ef0003	fffff	0	0	2007/03/31-19:50:14
sw172-22-46-224	1, 2	vsanid destid srcid	4001	ef0003	ea01a7	fffff	810	110.220K	2007/02/21-08:03:55
sw172-22-46-223	1, 4	vsanid destid srcid	1ef000f	650101	fffff	fffff	7.065K	296.600K	2007/04/02-11:43:05
sw172-22-46-174	10, 2	vsanid destid srcid	1ef000f	ee00ef	fffff	fffff	443	28.116K	2007/04/04-12:44:10
sw172-22-46-221	4, 3	vsanid destid srcid	1ee0006	6401da	fffff	fffff	631	37.440K	2007/04/04-12:44:05

Step 2 Click the **Create** icon to create a flow.



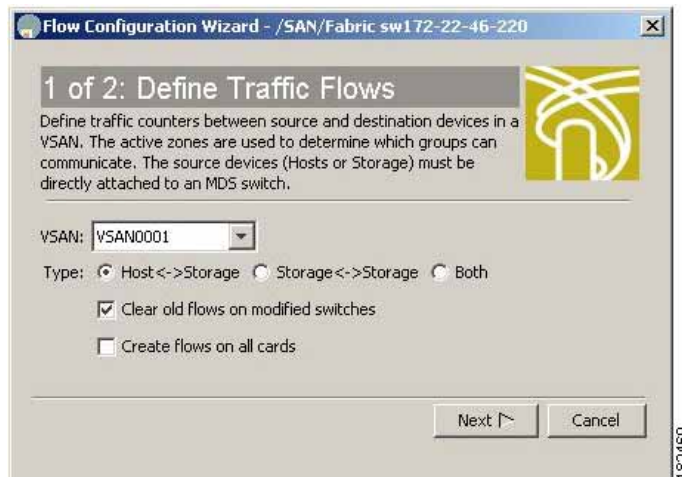
Note If you are managing your fabrics with Performance Manager, you need to set up an initial set of flows and collections on the fabric.



Note When creating flows, make sure the device type is set correctly.

You see the Define Traffic Flows dialog box as shown in Figure 29-3.

Figure 29-3 Define Traffic Flows Dialog Box



Step 3 Choose the VSAN from which to create flows. Flows are defined per active zone set.

Step 4 Click the **Type** radio button for the flow type you want to define.

Step 5 (Optional) Check the **Clear old flows on modified switches** check box if you want to remove old flow data.


Step 6 (Optional) Check the **Create flows on all cards** check box if you want to create flows on all cards for flexibility.

Step 7 Click **Next** to review the available flows for the chosen VSAN.

Send comments to nx5000-docfeedback@cisco.com

You see the Review Traffic Flows dialog box as shown in [Figure 29-4](#).

Figure 29-4 Review Traffic Flows Dialog Box



Step 8 Remove any flows you are not interested in.

Step 9 Click **Finish** to create the flow.



Nexus 5000 Management Software FAQ

This chapter answers some of the most frequently asked questions about Cisco Fabric Manager and Device Manager. This chapter contains the following sections:

- [Nexus 5000 Series Issues](#)
- [General Fabric Manager Issues](#)

Nexus 5000 Series Issues

This section includes issues that are specific to Cisco Nexus 5000 Series switches, and contains the following topics:

- [What is Display FCoE Mode?, page 30-1](#)
- [Switching to Display FCoE Mode, page 30-1](#)

What is Display FCoE Mode?

When Fabric Manager is running in Display FCoE mode, Fabric Manager displays additional tree nodes, menu items, toolbar buttons, and topology nodes/links related to Fibre Channel over Ethernet (FCoE).

These features include the following:

- Displaying Virtual Interface Group (VIG) wizard and Virtual Interface wizard toolbar buttons and menu items.
- Displaying VIG and virtual interface tree nodes in the physical attributes tree.
- Displaying FCoE links in the topology as dot-dot-dash lines.
- Displaying the Cisco Nexus 5000 Series switch with a different icon in the topology.

Switching to Display FCoE Mode

To switch Fabric Manager to Display FCoE mode, perform the following steps:

- Edit the `server.properties` file, which is located in the `<installation directory>/conf` folder of the Fabric Manager server.
- Search the file for the text “`displayFCoE`”, which is the display FCoE property.
- Change the `displayFCoE` property to **true**.

Send comments to nx5000-docfeedback@cisco.com

- Save the file and exit the editor.
- Restart the Fabric Manager server.

After you complete these steps, the file contains the following lines:

```
# FM Client and Server Attributes
# Display FCoE feature related tree nodes, menu items, toolbar buttons,
# topology nodes/links if Nexus 5k switches are present in the fabric
displayFCoE = true
```

General Fabric Manager Issues

For information about Fabric Manager frequently asked questions, see the Management Software FAQ available at the following location:

http://cisco.com/en/US/docs/storage/san_switches/mds9000/sw/rel_3_x/configuration/guides/fm_3_2/tsfm.html



Troubleshooting Your Fabric

This chapter describes basic troubleshooting methods used to resolve issues with switches. This chapter contains the following sections:

- [Troubleshooting Tools and Techniques, page 31-1](#)
- [Analyzing Switch Device Health, page 31-3](#)
- [Analyzing Switch Fabric Configuration, page 31-4](#)
- [Analyzing End-to-End Connectivity, page 31-5](#)
- [Using the Ping Tool \(fcping\), page 31-7](#)
- [Using Trace Route \(fctrace\) and Other Troubleshooting Tools, page 31-7](#)
- [Analyzing the Results of Merging Zones, page 31-8](#)
- [Using the Show Tech Support Command, page 31-9](#)
- [Running CLI Commands, page 31-11](#)
- [Locating Other Switches, page 31-12](#)
- [Fibre Channel Timeout Values, page 31-14](#)
- [Configuring a Fabric Analyzer, page 31-16](#)
- [Configuring World Wide Names, page 31-22](#)
- [Configuring a Secondary MAC Address, page 31-22](#)
- [FC ID Allocation for HBAs, page 31-23](#)
- [Default Settings, page 31-23](#)

Troubleshooting Tools and Techniques

Multiple techniques and tools are available to monitor and troubleshoot Cisco switches. These tools provide a complete, integrated, multilevel analysis solution.

Fabric Manager Server—The Cisco Fabric Manager Server provides a long-term, high-level view of storage network performance. Fabric-wide performance trends can be analyzed using Performance Manager. It provides the starting point for deeper analysis to resolve network hot-spots.

Device Manager—If a performance problem is detected with the Fabric Manager Server, use Cisco Device Manager to view port-level statistics in real-time. Details on protocols, errors, discards, byte and frame counts are available. Samples can be taken as frequently as every 2 seconds, and values can be viewed in text form or graphically as pie, bar, area, and line charts.

Send comments to nx5000-docfeedback@cisco.com

Traffic Analyzer—Another option is to launch the Cisco Traffic Analyzer for Fibre Channel from the Fabric Manager Server to analyze the traffic in greater depth. The Cisco Traffic Analyzer allows you to breakdown traffic by VSANs and protocols and to examine SCSI traffic at a logical unit number (LUN) level.

Protocol Analyzer—If even deeper investigation is needed, the Cisco Protocol Analyzer for Fibre Channel can be launched in-context from the Cisco Traffic Analyzer. The Cisco Protocol Analyzer enables you to examine actual sequences of Fibre Channel frames easily using the Fibre Channel and SCSI decoders Cisco developed for Wireshark.

Port Analyzer Adapter—Fabric Manager Server and Device Manager use SNMP to gather statistics. They fully utilize the built-in switch statistics counters.

Integration with the Cisco Traffic Analyzer and Cisco Protocol Analyzer extend the switch analysis capabilities by analyzing the Fibre Channel traffic itself. The Cisco Switched Port Analyzer (SPAN) enables these solutions using a flexible, nonintrusive technique to mirror traffic selectively from one or more ports to another switch port within a fabric.

The Cisco Port Analyzer Adapter (PAA) encapsulates SPAN traffic in an Ethernet header for transport to a PC or workstation for analysis. Both Fibre Channel control and data plane traffic are available using SPAN. The PAA broadcasts the Ethernet packets so they cannot be routed across IP networks. Hubs and switches can be used, provided they are in the same Ethernet subnet. Direct connections between a PAA and the PC are also supported. The PAA can reduce Ethernet traffic by truncating Fibre Channel data.

Both the Cisco Traffic Analyzer and Cisco Protocol Analyzer require the PAA to transport SPAN traffic to a PC or workstation.



Note

The Cisco Traffic Analyzer works best with the Cisco Port Analyzer Adapter 2, because it provides a length value for truncated data, enabling accurate byte count reporting.

Cisco Traffic Analyzer

The Cisco Traffic Analyzer for Fibre Channel provides real-time analysis of SPAN traffic or traffic captured previously using the Cisco Protocol Analyzer. The Fibre Channel traffic from multiple Cisco Port Analyzer Adapters (PAA) can be aggregated and analyzed by the Cisco Traffic Analyzer.

There are limits to how many SPAN sources can be sent to a single SPAN destination port. Aggregation extends the amount of information that can be analyzed in a unified set of reports by the Cisco Traffic Analyzer.



Note

The aggregation capabilities are restricted to the information collect by Ethernet connections to a single PC. Aggregation across multiple PCs is not available.

The Cisco Traffic Analyzer provides reports through a Web server, so you can view them locally or remotely. The traffic analysis functions are provided by “ntop” open-source software, which was enhanced by Cisco to add Fibre Channel and SCSI analysis and enhanced inter-switch link (ISL) header support for SPAN. ntop is available on the Cisco.com software download center, under the Cisco Port Analyzer Adapter. ntop is also available on the Internet at <http://www.ntop.org/ntop.html>. The Cisco enhanced ntop runs under Microsoft Windows and Linux operating systems.

Send comments to nx5000-docfeedback@cisco.com

The Cisco Traffic Analyzer for Fibre Channel presents reports with network wide statistics. The Summary Traffic report shows what percentage of traffic was within different ranges of frames sizes. A breakdown of the percentage of traffic for each protocol such as SCSI and ELS, is provided. The average and peak throughput for the SPAN traffic being analyzed are also provided.

Fibre Channel traffic can be analyzed on a per-VSAN basis with the Cisco Traffic Analyzer. The Domain Traffic Distribution graphs indicate how much traffic (bytes) were transmitted or received by a switch for a particular VSAN. FC Traffic Matrix graphs show how much traffic is transmitted and received between Fibre Channel sources and destinations. The total byte and frame counts for each VSAN are also provided.

Statistics can be analyzed for individual host and storage ports. You can see the percentage of SCSI read versus write traffic, SCSI versus other traffic, and percentage of transmitted versus received bytes and frames. The peak and average throughput values are available for data transmitted and received by each port.

Cisco Protocol Analyzer

The Cisco Protocol Analyzer for Fibre Channel enables you to view Fibre Channel traffic frames in real-time or from a capture file. Fibre Channel and SCSI decoders enable you to view and analyze traffic at the frame level. It matches response with request for complete decoding, which greatly simplifies navigation. Response time between response and status are presented.

The Cisco Protocol Analyzer is VSAN aware, so VSANs can be used as criteria for capture and display filters. VSAN numbers can also be displayed in a column. Summary statistics are available for protocol distribution percentages and total bytes or frames transferred between specific Fibre Channel source and destination pairs. File capture and filtering controls are available. Captured files can be analyzed by either the Cisco Protocol Analyzer or the Cisco Traffic Analyzer.

Numerous features have been included for ease-of-use. You can find frames that meet particular criteria and mark them. Entries in the frame (packet) list can be colored to highlight items of interest, and columns can be added or removed as desired.

The protocol analysis functions are provided by Wireshark open-source software, which was enhanced by Cisco to decode Fibre Channel and SCSI protocols and support enhanced inter-switch link (ISL) headers for SPAN. Wireshark is available on the Cisco.com software download center, under the Cisco Port Analyzer Adapter. Wireshark is also available on the Internet at <http://www.wireshark.org>. Wireshark runs under Microsoft Windows, Solaris, and Linux operating systems.

Analyzing Switch Device Health


The Switch Health option lets you determine the status of the components of a specific switch.

To use the Switch Health option in Fabric Manager to determine the status of the components of a specific switch, perform this task:

-
- Step 1** Choose **Tools > Switch Health**.
You see the Switch Health Analysis window.
- Step 2** Click **Start** to identify problems currently affecting the selected switch.
You see any problems listed in the switch health analysis window as shown in [Figure 31-1](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 31-1 Results of a Switch Health Analysis



Step 3 Click **Clear** to remove the contents of the Switch Health Analysis window.

Step 4 Click **Close** to close the window.

Analyzing Switch Fabric Configuration

The Fabric Configuration option lets you analyze the configuration of a switch by comparing the current configuration to a specific switch or to a policy file. You can save a switch configuration to a file, and then compare all switches against the configuration in the file.

To use the Fabric Configuration option in Fabric Manager to analyze the configuration of a switch, perform this task:

Step 1 Choose **Tools > Fabric Configuration**.

You see the Fabric Configuration Analysis dialog box.

Step 2 Decide whether you want to compare the selected switch to another switch, or to a policy file.

- If you are making a switch comparison, check **Policy Switch**, and then choose the drop-down arrow to see a list of switches.
- If you are making a policy comparison, check **Policy File**. Then click the ... button to the right of this option to browse your file system and choose a policy file (*.XML).

Step 3 Click **Rules** to set the rules to apply when running the Fabric Configuration Analysis tool.

You see the Rules window.

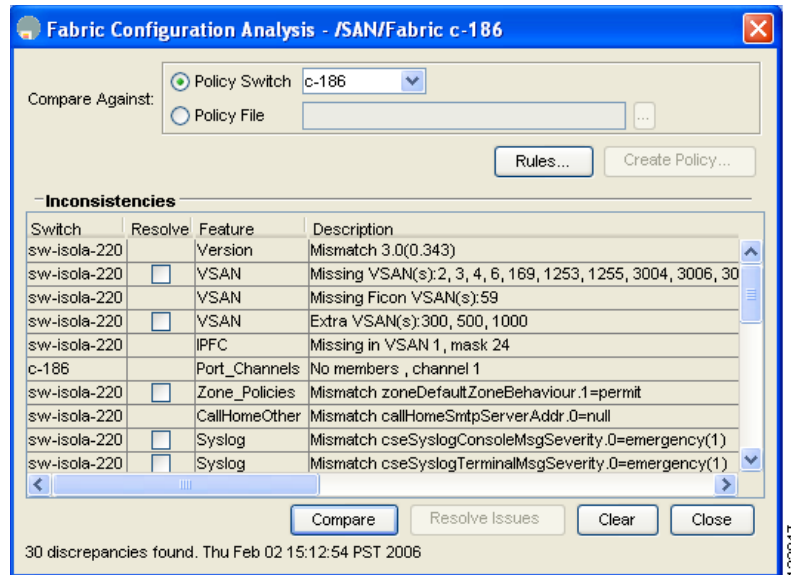
Step 4 Change the rules as needed and click **OK**.

Step 5 Click **Compare**.

The system analyzes the configuration and displays issues that arise as a result of the comparison as shown in [Figure 31-2](#).

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Figure 31-2 Results of a Fabric Configuration Analysis



- Step 6** Check the check boxes in the Resolve column for the issues you want to resolve.
- Step 7** To resolve the issues, click **Resolve Issues**.
- Step 8** Click **Clear** to remove the contents of the window.
- Step 9** Click **Close** to close the window.

Analyzing End-to-End Connectivity

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

To use the End to End Connectivity option in Fabric Manager to determine connectivity and routes, perform this task:

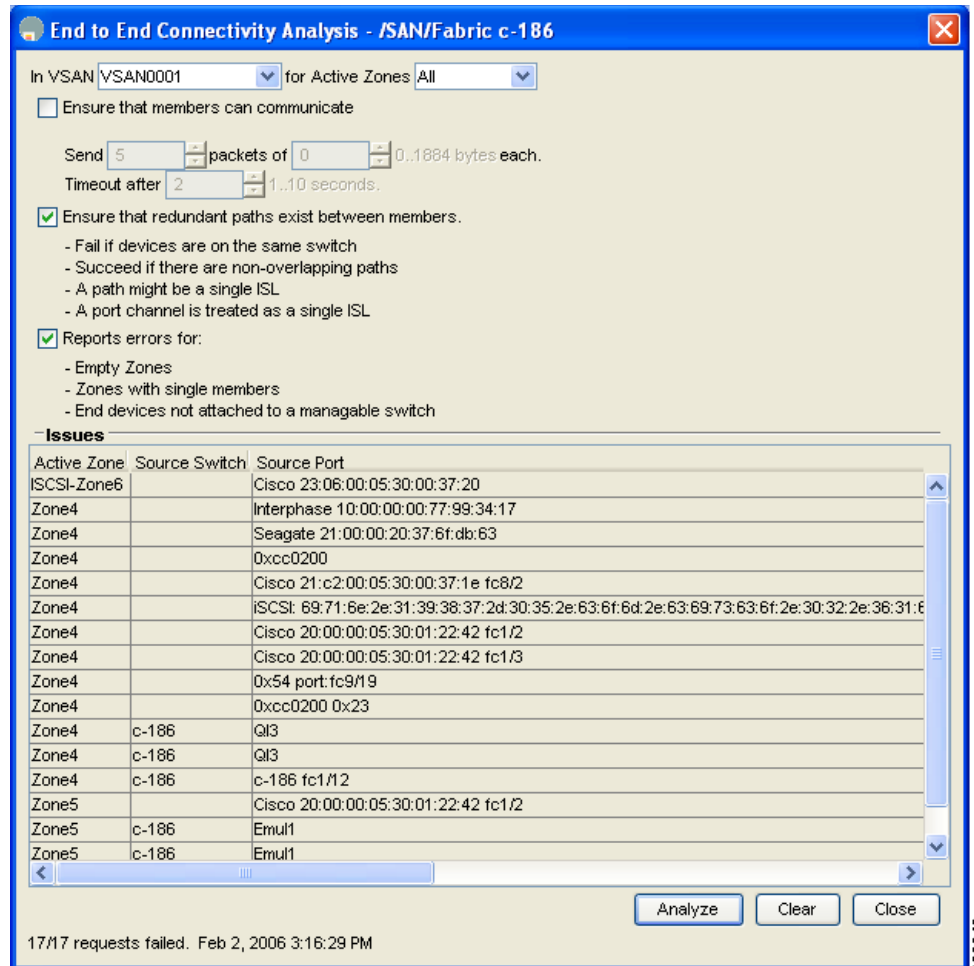
- Step 1** Choose **Tools > End to End Connectivity**.
You see the End to End Connectivity Analysis dialog box.
- Step 2** Choose the VSAN whose connectivity will be verified from the VSAN drop-down list.
- Step 3** Choose whether to perform the analysis for all active zones or for the default zone.
- Step 4** Click **Ensure that members can communicate** to perform a Fibre Channel ping between the selected endpoints.
- Step 5** Identify the number of packets, the size of each packet, and the timeout in milliseconds.
- Step 6** Analyze the redundant paths between endpoints by checking the **Ensure that redundant paths exist between members** check box.
- Step 7** Check the **Report errors for** check box to see a report of zone and device errors.

Send comments to nx5000-docfeedback@cisco.com

Step 8 Click **Analyze**.

The End to End Connectivity Analysis window displays the selected endpoints including the switch to which each is attached, and the source and target ports used to connect it, as shown in [Figure 31-3](#).

Figure 31-3 Results of an End-to-End Connectivity Analysis



The output shows all the requests that have failed:

- Ignoring empty zone—No requests are issued for this zone.
- Ignoring zone with single member—No requests are issued for this zone.
- Source/Target are unknown—No name server entries exist for the ports or we have not discovered the port during discovery.
- Both devices are on the same switch.
- No paths exist between the two devices.
- VSAN does not have an active zone set and the default zone is denied.
- Average time micro secs—The latency value was more than the threshold supplied.

Step 9 Click **Clear** to remove the contents of the window.

Send comments to nx5000-docfeedback@cisco.com

Step 10 Click **Close** to close the window.

Using the Ping Tool (fcping)

You can use the Ping tool to determine connectivity from another switch to a port on your switch.

To use the Ping tool in Fabric Manager to determine connectivity, perform this task:

Step 1 Choose **Tools > Ping**.

You can also select it from the right-click context menus for hosts and storage devices in the Fabric pane.

You see the Ping dialog box.

Step 2 Choose the source switch from the Source Switch drop-down list.


Step 3 Choose the VSAN in which you want to verify connectivity from the VSAN drop-down list.

Step 4 Choose the target end port for which to verify connectivity from the Target Endport drop-down list.

Step 5 Click **Start** to perform the ping between your switch and the selected port.

You see the Ping Results dialog box as shown in [Figure 31-4](#).

Figure 31-4 Ping Results



Step 6 Click **Clear** to clear the contents of the window and perform another ping, or click **Close** to close the window.

Using Trace Route (fctrace) and Other Troubleshooting Tools

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

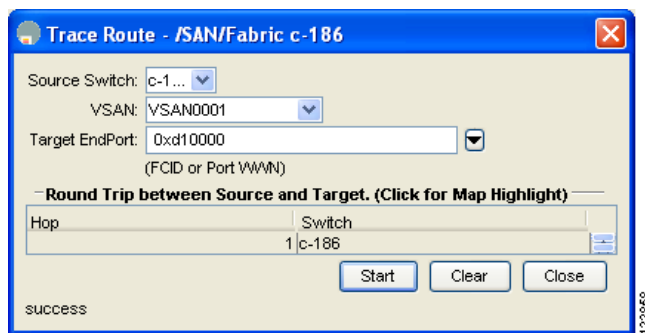
- Trace Route—Verify connectivity between two end devices that are currently selected on the Fabric pane.
- Device Manager—Launch the Device Manager for the switch selected on the Fabric pane.
- Command Line Interface—Open a Telnet or SSH session for the switch selected on the Fabric pane.

Send comments to nx5000-docfeedback@cisco.com

To use the Trace Route option in Fabric Manager to verify connectivity, perform this task:

-
- Step 1** Choose **Tools > Trace Route**.
You see the Trace Route dialog box.
 - Step 2** Choose the source switch from the Source Switch drop-down list.
 - Step 3** Choose the VSAN for which to verify connectivity from the VSAN drop-down list.
 - Step 4** Choose the target end port for which to verify connectivity from the Target Endport drop-down list.
 - Step 5** Click **Start** to perform the traceroute between your switch and the selected port.
You see the results at the bottom of the dialog box as shown in [Figure 31-5](#).

Figure 31-5 Successful Trace Route Results



- Step 6** Click **Clear** to clear the contents of the window and perform another traceroute, or click **Close** to close the window.
-

Analyzing the Results of Merging Zones

You can use the Zone Merge option on the Zone menu to determine if two connected switches have compatible zone configurations.

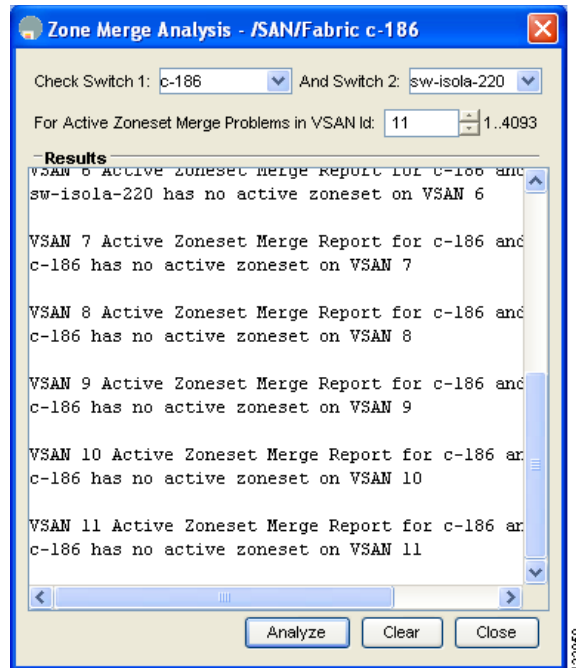
To use the Zone Merge option in Fabric Manager to determine zone configuration compatibility, perform this task:

-
- Step 1** Choose **Zone > Merge Analysis**.
You see the Zone Merge Analysis dialog box.
 - Step 2** Choose a switch from each drop-down list.
 - Step 3** Choose the VSAN for which you want to perform the zone merge analysis.
 - Step 4** Repeat Step 3 as needed.
 - Step 5** Click **Analyze**.

The Zone Merge Analysis window displays any inconsistencies between the zone configuration of the two selected switches as shown in [Figure 31-6](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 31-6 Results of Zone Merge Analysis



Step 6 Click **Clear** to remove the contents of the window.

Step 7 Click **Close** to close the window.

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 can be provided to technical support representatives when reporting a problem.

You can enter the **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** command output and the map file image) and send the resulting zipped file to technical support.

To enter the **show tech support** command using Fabric Manager, perform this task:

Step 1 Choose **Tools > Show Tech Support**.

You see the Show Tech Support dialog box.

Step 2 Choose the switches for which to view technical support information by checking the check boxes for each switch.

Step 3 Set the time-out value.

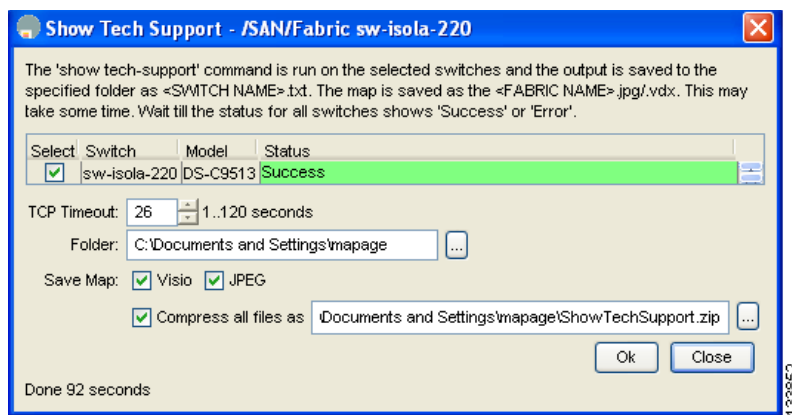
Send comments to nx5000-docfeedback@cisco.com

The default is 30 seconds.

- Step 4** Select the folder where you want the text files (containing the **show tech support** command information) to be written.
- Step 5** Check the **Save Map** check box if you want to save a screenshot of your map as a JPG file.
- Step 6** Check the **Compress all files as** check box to compress the files into a zip file.
- Step 7** Click **OK** to start the **show tech support** command on the switches that you specified, or click **Close** to close the Show Tech Support dialog box without using the **show tech support** command (see [Figure 31-7](#)).

In the Status column next to each switch, you see a highlighted status. A yellow highlight indicates that the **show tech support** command is currently running on that switch. A red highlight indicates an error. A green highlight like the one shown in [Figure 31-7](#) indicates that the **show tech support** command has completed successfully.

Figure 31-7 Successful Results of the **show tech support** Command



- Step 8** If prompted, enter your user name and password in the appropriate fields for the switch in question.



Note In order for Fabric Manager to successfully enter the **show tech support** command on a switch, that switch must have this user name and password. Fabric Manager is unable to log into a switch that does not have a user name and password and an error is returned for that switch.



Note If you would like to view output files of the **show tech support** command without using Fabric Manager, open them with any text editor. Each file is named with the switch's IP address and has a .TXT extension (for example, 111.22.33.444.txt).

Send comments to nx5000-docfeedback@cisco.com

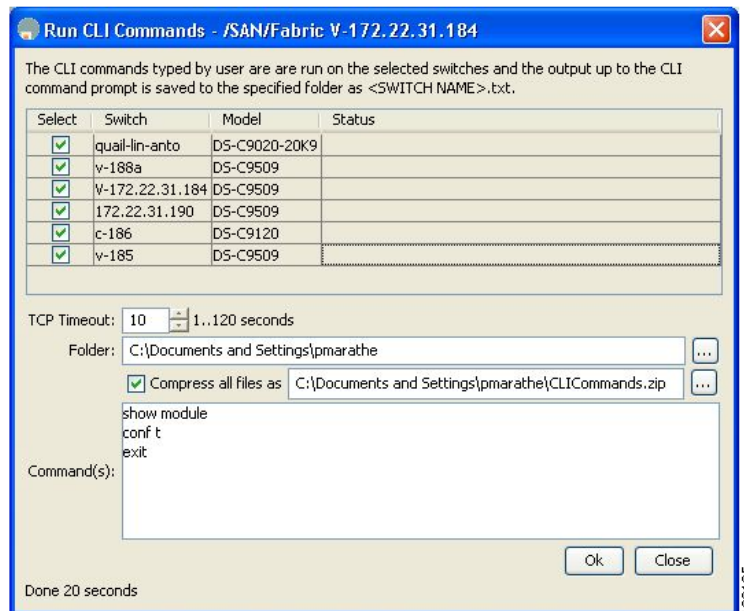
Running CLI Commands

You can use the Run CLI Commands feature to run a CLI command on multiple switches. To run CLI commands using Fabric Manager, perform this task:

Step 1 Choose **Tools > Run CLI Commands**.

You see the Run CLI Commands dialog box with all switches selected as shown in [Figure 31-8](#).

Figure 31-8 Run CLI Commands Dialog Box



Step 2 Uncheck the check box for the switch(es) for which you do not want to run CLI commands.

Step 3 Specify where you want the file to be saved.



Note A separate report is issued for each switch. Check the reports to verify whether a CLI command failed.

Step 4 Enter the command(s) in the Command(s) text box. If the commands are configuration mode commands, you must also enter the **exit** command.




Note For the commands to execute, you cannot be in configuration mode.

Step 5 Click **OK** to run the CLI command(s).

You see the Run CLI Commands dialog box showing the status of each switch as shown in [Figure 31-9](#).

Send comments to nx5000-docfeedback@cisco.com

Figure 31-9 Run CLI Commands Status



Step 6 Click **Close** to close the dialog box.

Adjusting for Daylight Savings Time



Note

Starting in 2007, daylight savings time in the United States starts on the second Sunday in March and ends on the first Sunday in November.

You can use the Run CLI Commands feature in Fabric Manager to adjust the time change configuration in your switches. Enter the following commands in the Command(s) text box:

```
configure
no clock summer-time
clock summer-time daylight_timezone_name 2 Sunday March 02:00 1 Sunday November 02:00 60
```

Locating Other Switches

The Locate Switches option uses SNMPv2 and discovers devices responding to SNMP requests with the read-only community string public. You can use this feature in these situations:

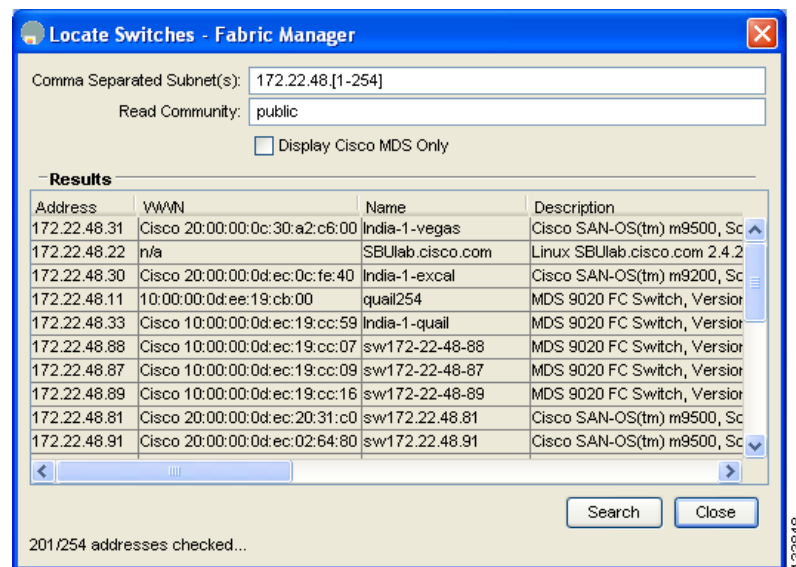
- You have third-party switches that do not implement the FC-GS3 FCS standard that provides management IP addresses.
- You want to locate other Cisco switches in the subnet but are not physically connected to the fabric.

Send comments to nx5000-docfeedback@cisco.com

To locate switches that are not included in the currently discovered fabric using Fabric Manager, perform this task:

- Step 1** Choose **File > Locate Switches and Devices**.
- You see the Locate Switches dialog box.
- Step 2** In the Comma Separated Subnets field, enter a range of specific addresses belonging to a specific subnet to limit the research for the switches.
- To look for a Cisco 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 Nexus 5000 family Only** to display only the Cisco Nexus 5000 Series switches in your network fabric.
- Step 5** Click **Search** to discover switches and devices in your network fabric.
- You see the results of the discovery in the Locate Switches window. (See [Figure 31-10](#).)

Figure 31-10 Search Results for Switches and Devices



Note The number in the lower left corner of the screen increments as the device locator attempts to discover the devices in your network fabric. When the discovery process is complete, the number indicates the number of rows displayed.

Send comments to nx5000-docfeedback@cisco.com

Step 6 Click **Close** to close this dialog box.

Fibre Channel Timeout Values

You can modify Fibre Channel protocol related timer values for the switch by configuring the following timeout values (TOVs):

- Distributed services TOV (D_S_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 5,000 milliseconds.
- Error detect TOV (E_D_TOV)—The valid range is from 1,000 to 10,000 milliseconds. The default is 2,000 milliseconds. This value is matched with the other end during port initialization.
- Resource allocation TOV (R_A_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 10,000 milliseconds. This value is matched with the other end during port initialization.



Note

The fabric stability TOV (F_S_TOV) constant cannot be configured.



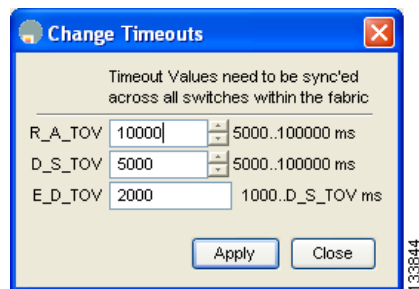
Caution

The D_S_TOV, E_D_TOV, and R_A_TOV values cannot be globally changed unless all VSANs in the switch are suspended.

To configure timeouts using Fabric Manager, perform this task:

- Step 1** Choose **SAN** in the Logical Domains pane to include all VSANs.
- Step 2** Expand **Switches**, expand **FC Services**, and choose **Timers & Policies** in the Physical Attributes pane. You see the timers for switches in the Information pane.
- Step 3** Click **Change Timeouts** to configure the timeout values. You see the Change Timeouts dialog box as shown in [Figure 31-11](#).

Figure 31-11 *Change Timeouts Dialog Box*



- Step 4** Indicate values for R_A_TOV (Resource Allocation Timeout Value), D_S_TOV (Distributed Services Timeout Value), and E_D_TOV (Error Detect Timeout Value).
- Step 5** Click **Apply**.

Send comments to nx5000-docfeedback@cisco.com


Step 6 Click **Close** to close the dialog box.

To configure timer policies in Device Manager, perform this task:

Step 1 Choose **FC > Advanced > Timers/Policies**.

You see timer policies for a single switch in the dialog box as shown in [Figure 31-12](#).

Figure 31-12 Configure Timer Policies in Device Manager



Step 2 Choose a network from the drop-down list and specify a switch.

Step 3 Check the check boxes for **InOrderDeliver** and/or **Trunk Protocol**.

Step 4 Click **Apply**.

Step 5 Click **Close** to close the dialog box.

Timer Configuration Per-VSAN

You can also issue an `ftimer` for a specified VSAN to configure different TOV values for VSANs with special links such as FC or IP tunnels. You can configure different `E_D_TOV`, `R_A_TOV`, and `D_S_TOV` values for individual VSANs. Active VSANs are suspended and activated when their timer values are changed.



Caution

You cannot perform a nondisruptive downgrade to any earlier version that does not support per-VSAN FC timers.



Note

This configuration must be propagated to all switches in the fabric; be sure to configure the same value in all switches in the fabric.

To configure per-VSAN FC timers using Fabric Manager, perform this task:

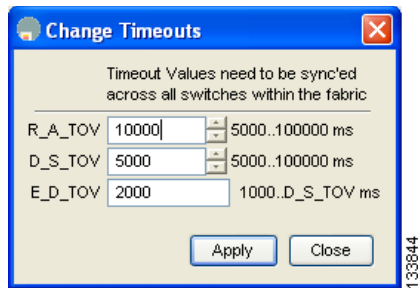
Step 1 Choose the VSAN for timer configuration from the Logical Domains pane.

Send comments to nx5000-docfeedback@cisco.com

If a VSAN is not specified when you change the policies, the changed value is applied to all VSANs in the switch.

- Step 2** Expand **Switches**, expand **FC Services**, and choose **Timers & Policies** in the Physical Attributes tree. You see timeouts for only switches in the selected VSAN shown in the Information pane.
- Step 3** Click **Change Timeouts** to configure the time-out values. You see the Change Timeouts dialog as shown in [Figure 31-13](#).

Figure 31-13 Change Timeouts per VSAN in Fabric Manager



- Step 4** Change the timeout values shown in [Figure 31-13](#).
- Step 5** Indicate values for R_A_TOV (Resource Allocation Timeout Value), D_S_TOV (Distributed Services Timeout Value), and E_D_TOV (Error Detect Timeout Value).
- Step 6** Click **Apply**.
- Step 7** Click **Close** to close the dialog box.

Configuring a Fabric Analyzer

Fibre Channel protocol analyzers capture, decode, and analyze frames and ordered sets on a link. Existing Fibre Channel analyzers can capture traffic at wire rate speed. They are expensive and support limited frame decoding. Also, existing analyzers disrupt the traffic on the link while snooping traffic.

With the Cisco Fabric Analyzer, you can capture Fibre Channel control traffic from a switch and decode it without having to disrupt any connectivity, and without having to be local to the point of analysis.

The Cisco Fibre Channel protocol analyzer is based on two popular public-domain software applications:

- libpcap—See <http://www.tcpdump.org>.
- Wireshark—See <http://www.wireshark.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.

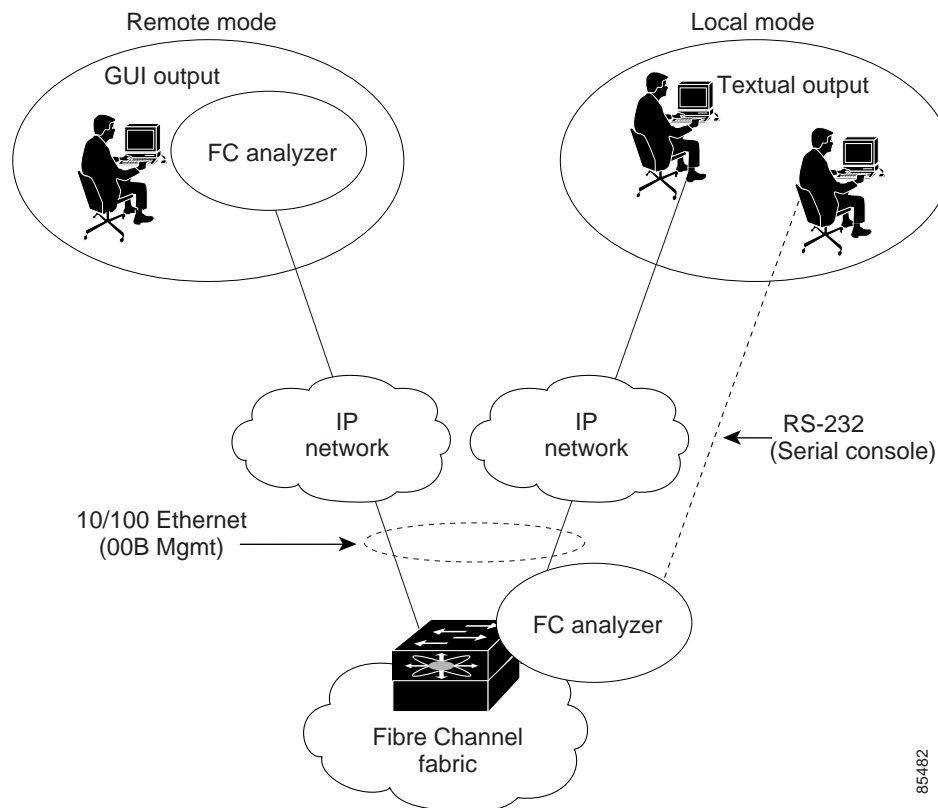
[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

About the Cisco Fabric Analyzer

The Cisco Fabric Analyzer consists of two separate components (see [Figure 31-14](#)):

- Software that runs on the Cisco Cisco Nexus 5000 Series switch and supports two modes of capture:
 - A text-based analyzer that supports local capture and decodes captured frames
 - A daemon that supports remote capture
- GUI-based client that runs on a host that supports libpcap such as Windows or Linux and communicates with the remote capture daemon in a Cisco Cisco Nexus 5000 Series switch.

Figure 31-14 Cisco Fabric Analyzer Use



85482

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 Cisco Nexus 5000 Series 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 Cisco Nexus 5000 Series switch, it is protected by the roles-based policy that limits access in each switch.

Send comments to nx5000-docfeedback@cisco.com

Remote Capture Daemon

This daemon is the server end of the remote capture component. The Wireshark analyzer running on a host is the client end. They communicate with each other using the Remote Capture Protocol (RPCAP). RPCAP uses two endpoints, a TCP-based control connection and a TCP or UDP-based data connection based on TCP (default) or UDP. The control connection is used to remotely control the captures (start or stop the capture, or specify capture filters). Remote capture can only be performed to explicitly configured hosts. This technique prevents an unauthorized machine in the network from snooping on the control traffic in the network.

RPCAP supports two setup connection modes based on firewall restrictions:

- **Passive mode (default)**—The configured host initiates connection to the switch. Multiple hosts can be configured to be in passive mode and multiple hosts can be connected and receive remote captures at the same time.
- **Active mode**—The switch initiates the connection to a configured host, one host at a time.

Using capture filters, you can limit the amount of traffic that is actually sent to the client. Capture filters are specified at the client end on Wireshark, not on the switch.

GUI-Based Client

The Wireshark 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.wireshark.org>. The Wireshark GUI front-end supports a rich interface such as a colored display, graphical help in defining filters, and specific frame searches. These features are documented on Wireshark's website.

While remote capture through Wireshark supports capturing and decoding Fibre Channel frames from a Cisco Cisco Nexus 5000 Series switch, the host running Wireshark does not require a Fibre Channel connection to the switch. The remote capture daemon running on the switch sends the captured frames over the out-of-band Ethernet management port. This capability allows you to capture and decode Fibre Channel frames from your desktop or laptop.

Configuring the Cisco Fabric Analyzer

You can configure the Cisco Fabric Analyzer to perform one of two captures:

- **Local capture**—A local capture cannot be saved to persistent storage or synchronized to standby. It launches the textual version on the fabric analyzer directly on the console screen. The capture can also be saved on the local file system.
- **Remote capture**—A remote capture can be saved to persistent storage. It can be synchronized to the standby supervisor module and a stateless restart can be issued, if required.

To use the Cisco Fabric Analyzer feature, traffic should be flowing to or from the supervisor module.

Send comments to nx5000-docfeedback@cisco.com

Sending Captures to Remote IP Addresses



Caution

You must use the eth2 interface to capture control traffic on a supervisor module.

To capture remote traffic, use one of the following options:

- The capture interface can be specified in Wireshark as the remote device:

```
rpcap://<ipaddress or switch hostname>/eth2
```

For example:

```
rpcap://cp-16/eth2
rpcap://17.2.1.1/eth2
```

- The capture interface can be specified either in the capture dialog box or by using the `-i` option at the command line when invoking Wireshark.

```
wireshark -i rpcap://<ipaddress|hostname>[:<port>]/<interface>
```

For example:

```
wireshark -i rpcap://172.22.1.1/eth2
```

or

```
wireshark -i rpcap://customer-switch.customer.com/eth2
```



Note

For example, in a Windows 2000 setup, click **Start** on your desktop and choose **Run**. In the resulting Run window, type the required command line option in the Open field.

Displaying Captured Frames

You can selectively view captured frames by using the display filters feature. For example, instead of viewing all the frames from a capture, you may only want to view Exchange Link Protocol (ELP) request frames. This feature only limits the captured view; it does not affect the captured or the saved frames. Procedures to specify, use, and save display filters are already documented in the Wireshark website (<http://www.wireshark.org>).

These examples show how to use this feature:

- To view all packets in a specified VSAN, use this expression:

```
mdshdr.vsan == 2
```

- To view all SW_ILS frames, use this expression:

```
fcswils
```

- To view class F frames, use this expression:

```
mdshdr.sof == SOFf
```

- To view all FSPF frames, use this expression:

```
swils.opcode == HLO || swils.opcode == LSU || swils.opcode == LSA
```

Send comments to nx5000-docfeedback@cisco.com

- To view all FLOGI frames, use this expression:
`fcels.opcode == FLOGI`
- To view all FLOGI frames in VSAN 1, use this expression:
`fcels.opcode == FLOGI && mdshdr.vsan == 2`
- To view all name server frames, use this expression:
`dNS`

Defining Display Filters

Display filters limit the frames that can be displayed, but not what is captured (similar to any view command). The filters to be displayed can be defined in multiple ways in the GUI application:

- Auto-definition
- Manual definition
- Assisted manual definition
- Only manual definition in local capture
- No assists

Regardless of the definition, each filter must be saved and identified with a name.



Note

This GUI-assisted feature is part of Wireshark, and you can obtain more information from <http://www.wireshark.org>.

Capture Filters

You can limit what frames are captured by using the capture filters feature in a remote capture. This feature limits the frames that are captured and sent from the remote switch to the host. For example, you can capture only class F frames. Capture filters are useful in restricting the amount of bandwidth consumed by the remote capture.

Unlike display filters, capture filters restrict a capture to the specified frames. No other frames are visible until you specify a completely new capture.

The syntax for capture filter is different from the syntax for display filters. Capture filters use the Berkeley Packet Filter (BPF) library that is used in conjunction with the libpcap freeware. The list of all valid Fibre Channel capture filter fields are provided later in this section.

Procedures to configure capture filters are already documented in the Wireshark website (<http://www.wireshark.org>).

These examples show how to use this feature:

- To capture frames only on a specified VSAN, use this expression:
`vsan = 1`
- To capture only class F frames, use this expression:
`class_f`

Send comments to nx5000-docfeedback@cisco.com

- To capture only class Fibre Channel ELS frames, use this expression:

```
els
```

- To capture only name server frames, use this expression:

```
dns
```

- To capture only SCSI command frames, use this expression:

```
fcp_cmd
```



Note

This feature is part of libpcap and you can obtain more information from <http://www.tcpdump.org>.

Permitted Capture Filters

This section lists the permitted capture filters:

```
o vsan
o src_port_idx
o dst_port_idx
o sof
o r_ctl
o d_id
o s_id
o type
o seq_id
o seq_cnt
o ox_id
o rx_id
o els
o swils
o fcp_cmd (FCP Command frames only)
o fcp_data (FCP data frames only)
o fcp_rsp (FCP response frames only)
o class_f
o bad_fc
o els_cmd
o swils_cmd
o fcp_lun
o fcp_task_mgmt
o fcp_scsi_cmd
o fcp_status
o gs_type (Generic Services type)
o gs_subtype (Generic Services subtype)
o gs_cmd
o gs_reason
o gs_reason_expl
o dns (name server)
o udns (unzoned name server)
o fcs (fabric configuration server)
o zs (zone server)
o fc (use as fc[x:y] where x is offset and y is length to compare)
o els (use as els[x:y] similar to fc)
o swils (use as swils[x:y] similar to fc)
o fcp (use as fcp[x:y] similar to fc)
o fcct (use as fcct[x:y] similar to fc)
```

Send comments to nx5000-docfeedback@cisco.com

Configuring World Wide Names

The world wide name (WWN) in the switch is equivalent to the Ethernet MAC address. As with the MAC address, you must uniquely associate the WWN to a single device. The principal switch selection and the allocation of domain IDs rely on the WWN. The WWN manager, a process-level manager residing on the switch's supervisor module, assigns WWNs to each switch.

Cisco Cisco Nexus 5000 Series switches support three network address authority (NAA) address formats (see [Table 31-1](#)).

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 world-wide names should be made by an administrator or individual who is completely familiar with switch operations.

Link Initialization WWN Usage

Exchange Link Protocol (ELP) and Exchange Fabric Protocol (EFP) use WWNs during link initialization. The usage details differ based on the Cisco SAN-OS software release:

Both ELPs and EFPs use the VSAN WWN by default during link initialization. However, the ELP usage changes based on the peer switch's usage:

- If the peer switch ELP uses the switch WWN, then the local switch also uses the switch WWN.
- If the peer switch ELP uses the VSAN WWN, then the local switch also uses the VSAN WWN.


Configuring a Secondary MAC Address

To allocate a secondary MAC address, perform this task:

-
- Step 1** Select a SAN (or a VSAN) from the Logical Domains pane.
You see a list of switches in the Information pane.
 - Step 2** Expand **Switches**, expand **FC Services**, and choose **WWN Manager** in the Physical Attributes pane.
 - Step 3** In the Information pane, scroll until you see the switch on which you want to configure a secondary MAC address (see [Figure 31-15](#)).

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Figure 31-15 Setting Secondary MAC Addresses



- Step 4 Enter the secondary MAC address in the **Secondary Mac Base** field.
- Step 5 Enter the range for the secondary MAC address in the **Secondary Mac Range** field.
- Step 6 Click the **Apply Changes** icon.

Displaying WWN Information

To display the status of the WWN configuration, perform this task:

- Step 1 Select a SAN (or a VSAN) from the Logical Domains pane.
You see a list of switches in the Information pane.
- Step 2 Choose **Switches > FC Services > WWN Manager** from the Physical Attributes pane.
You see the WWN information for each switch in the SAN or VSAN.

FC ID Allocation for HBAs

Fibre Channel standards require a unique FC ID to be allocated to an N port attached to a Fx port in any switch. To conserve the number of FC IDs used, Cisco Cisco Nexus 5000 Series switches use a special allocation scheme. See the [“FC ID Allocation for HBAs”](#) section on page 31-23.

Default Settings

Table 31-2 lists the default settings for the features included in this chapter.

Table 31-2 Default Settings for Advanced Features

Parameters	Default
CIM server	Disabled
CIM server security protocol	HTTP
D_S_TOV	5,000 milliseconds

[Send comments to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Table 31-2 *Default Settings for Advanced Features (continued)*

Parameters	Default
E_D_TOV	2,000 milliseconds
R_A_TOV	10,000 milliseconds
Time-out period to invoke fctrace	5 seconds
Number of frame sent by the fcping feature	5 frames
Remote capture connection protocol	TCP
Remote capture connection mode	Passive
Local capture frame limit s	10 frames
FC ID allocation mode	Auto mode
Loop monitoring	Disabled



I N D E X

Symbols

* (asterisk)

- autolearned entries [24-18](#)
- port security wildcards [24-13](#)

A

AAA

- DHCHAP authentication [23-9](#)

active zone sets

- considerations [16-4](#)
- enabling distribution [16-21](#)

administrative states

- description [10-5](#)
- setting [10-9](#)

authentication

- fabric security [23-1](#)

auto port mode

- description [10-4](#)

B

BB_credits

- configuring [10-12](#)
- description [10-6](#)
- reason codes [10-6](#)

bit errors

- reasons [10-11](#)

Brocade

- native interop mode [22-8](#)

buffer-to-buffer credits. See [BB_credits](#)

build fabric frames

- description [11-3](#)

C

company IDs

- FC ID allocations [22-6](#)

configuring NPV [12-4](#)

D

dead time intervals

- description [18-7](#)

default zones

- configuring [16-16](#)
- configuring access permissions [16-16](#)
- configuring policies [16-12](#)
- description [16-16](#)
- interoperability [22-9](#)
- policies [16-16](#)

destination IDs

- exchange based [14-3](#)
- flow based [14-3](#)
- in-order delivery [18-12](#)
- path selection [15-11](#)

device alias databases

- committing changes [17-6](#)
- discarding changes [17-7](#)
- locking the fabric [17-5](#)
- merging [17-8](#)

device aliases

- comparison with zones (table) [17-2](#)
- creating (procedure) [17-5](#)
- default settings [17-9](#)

Send comments to nx5000-docfeedback@cisco.com

- description [17-1](#)
- displaying information [17-8 to ??](#)
- enhanced mode [17-3](#)
- features [17-1](#)
- import legacy zone aliases [17-7](#)
- modifying databases [17-2](#)
- requirements [17-2](#)
- using [17-7](#)
- zone alias conversion [17-7](#)

DHCHAP

- AAA authentication [23-9](#)
- authentication modes [23-4](#)
- configuring [23-2 to ??](#)
- configuring AAA authentication [23-9](#)
- default settings [23-10](#)
- displaying security information [23-9](#)
- enabling [23-3, 23-4](#)
- group settings [23-6](#)
- hash algorithms [23-5](#)
- passwords for local switches [23-7](#)
- passwords for remote devices [23-8](#)
- timeout values [23-8](#)
- See also FC-SP [23-1](#)

Diffie-Hellman Challenge Handshake Authentication Protocol. See DHCHAP

documentation

- additional publications [1-iv](#)
- related documents [1-iv](#)

domain IDs

- allowed lists [11-10](#)
- assignment failures [10-7](#)
- configuring allowed lists [11-11](#)
- configuring CFS distribution [11-11](#)
- configuring fcalias members [16-17](#)
- description [11-8](#)
- distributing [11-1](#)
- enabling contiguous assignments [11-14](#)
- interoperability [22-8](#)
- preferred [11-10](#)

- static [11-10](#)

domain manager

- fast restart feature [11-4](#)
- isolation [10-7](#)

drop latency time

- configuring [18-15](#)

E

EFMD

- fabric binding [25-1](#)
- fabric binding initiation [25-3](#)

EISLs

- port channel links [14-1](#)

enhanced zones

- changing from basic zones [16-29](#)
- description [16-29](#)
- enabling [16-31](#)
- modifying database [16-31](#)

E port mode

- classes of service [10-3](#)
- description [10-3](#)

E ports

- fabric binding checking [25-2](#)
- FCS support [26-1](#)
- FSPF topologies [18-1](#)
- isolation [10-7](#)
- recovering from link isolations [16-23](#)
- trunking configuration [13-3](#)

Exchange Fabric Membership Data. See EFMD

exchange IDs

- in-order delivery [18-12](#)
- path selection [15-11](#)

exchange link parameter. See ELP

expansion port mode. See E port mode

Send comments to nx5000-docfeedback@cisco.com

F

fabric binding

- activation [25-5](#)
- checking for E ports [25-2](#)
- checking for TE ports [25-2](#)
- clearing statistics [25-9](#)
- compatibility with DHCPAP [23-3](#)
- copying to config database [25-6](#)
- copying to configuration file (procedure) [25-8](#)
- creating config database (procedure) [25-7](#)
- default settings [25-10](#)
- deleting databases [25-9](#)
- deleting from config database (procedure) [25-7](#)
- description [25-1](#)
- EFMD [25-1](#)
- enabling [25-3](#)
- enforcement [25-2](#)
- forceful activation [25-6](#)
- initiation process [25-3](#)
- licensing requirements [25-1](#)
- port security comparison [25-1](#)
- saving to config database [25-6](#)
- sWWN lists [25-4](#)
- viewing active databases (procedure) [25-8](#)
- viewing EFMD statistics (procedure) [25-8](#)
- viewing violations (procedure) [25-8](#)

Fabric Configuration Servers. See FCSs

Fabric-Device Management Interface. See FDMI

fabric login. See FLOGI

fabric port mode. See F port mode

fabric pWWNs

- zone membership [16-2](#)

fabric reconfiguration

- fcdomain phase [11-1](#)

fabrics

- See also build fabric frames

fabrics. See RCFs; build fabric frames [11-3](#)

fabric security

authentication [23-1](#)

default settings [23-10](#)

Fabric Shortest Path First. See FSPF

fabric WWNs. See fWWNs

fault tolerant fabrics

example (figure) [18-2](#)

fcaliases

adding members [16-18](#)

cloning [16-27](#)

configuring for zones [16-17](#)

creating [16-17](#)

renaming [16-26](#)

using [17-7](#)

fcdomains

autoreconfigured merged fabrics [11-7](#)

configuring CFS distribution [11-11](#)

default settings [11-21](#)

description [11-1](#)

displaying information [11-20](#)

domain IDs [11-8](#)

domain manager fast restart [11-4](#)

enabling autoreconfiguration [11-7](#)

incoming RCFs [11-6](#)

initiation [11-5](#)

overlap isolation [10-7](#)

restarts [11-2](#)

FC IDs

allocating [11-1, 22-6](#)

allocating default company ID lists [22-6](#)

allocation for HBAs [22-6](#)

configuring fcalias members [16-17](#)

description [11-15](#)

persistent [11-15, 11-20](#)

FC-SP

authentication [23-1](#)

enabling [23-4](#)

enabling on ISLs [23-9](#)

See also DHCPAP [23-1](#)

FCSs

Send comments to nx5000-docfeedback@cisco.com

- characteristics [26-2](#)
- configuring names [26-2](#)
- default settings [26-6](#)
- description [26-1](#)
- displaying fabric ports using Device Manager [26-5](#)
- displaying information [26-3 to ??](#)
- fctimers
 - displaying configured values [22-5](#)
 - distribution [22-4](#)
- FDMI
 - displaying database information [19-4](#)
- Fibre Channel
 - sWWNs for fabric binding [25-4](#)
 - timeout values [22-1 to ??](#)
- Fibre Channel domains. See [fcdomains](#)
- Fibre Channel interfaces
 - administrative states [10-5](#)
 - BB_credits [10-6](#)
 - configuring [10-8](#)
 - configuring descriptions [10-9, 10-10](#)
 - configuring frame encapsulation [10-10](#)
 - deleting from port channels [14-14](#)
 - disabling [10-9](#)
 - displaying information [10-14](#)
 - enabling [10-9](#)
 - operational states [10-5](#)
 - reason codes [10-5](#)
 - states [10-4](#)
 - See also [interfaces](#) [10-4](#)
- Fibre Channel Security Protocol. See [FC-SP](#)
- FLOGI
 - description [19-1](#)
- F port mode
 - classes of service [10-4](#)
 - description [10-3](#)
- F ports
 - description [10-3](#)
 - See also [Fx ports](#)
- frame encapsulation
 - configuring [10-10](#)
- FSPF
 - clearing VSAN counters [18-5](#)
 - computing link cost [18-6](#)
 - configuring globally [18-3 to ??](#)
 - configuring Hello time intervals [18-6](#)
 - configuring on a VSAN [18-4](#)
 - configuring on interfaces [18-5 to ??](#)
 - dead time intervals [18-7](#)
 - default settings [18-16](#)
 - description [18-1](#)
 - disabling [18-5](#)
 - disabling on interfaces [18-8](#)
 - disabling routing protocols [18-5](#)
 - enabling [18-5](#)
 - fault tolerant fabrics [18-2](#)
 - in-order delivery [18-12](#)
 - interoperability [22-9](#)
 - link state record defaults [18-3](#)
 - reconvergence times [18-2](#)
 - redundant links [18-2](#)
 - resetting configuration [18-5](#)
 - resetting to defaults [18-4](#)
 - retransmitting intervals [18-7](#)
 - routing services [18-1](#)
 - topology examples [18-2](#)
- FSPF routes
 - configuring [18-11](#)
 - description [18-11](#)
- full zone sets
 - considerations [16-4](#)
 - enabling distribution [16-21](#)
- fWWNs
 - configuring fcalias members [16-17](#)
- Fx ports
 - VSAN membership [15-4](#)

Send comments to nx5000-docfeedback@cisco.com

H

hard zoning

description [16-21](#)

HBA ports

configuring area FCIDs [11-17](#)

HBAs

FC ID allocations [22-6](#)

Hello time intervals

configuring for FSPF [18-6](#)

description [18-6](#)

I

indirect link failures

recovering [27-1](#)

in-order delivery

displaying status [18-15](#)

enabling globally [18-14](#)

guidelines [18-14](#)

reordering port channel frames [18-13](#)

interfaces

adding to port channels [14-12](#), [14-13](#)

assigning to VSANs [15-8](#)

configuring data field size [10-11](#)

configuring descriptions [10-9](#), [10-10](#)

configuring fcalias members [16-17](#)

default settings [10-14](#)

deleting from port channels [14-14](#)

displaying information [10-14](#)

forced addition to port channels [14-13](#)

isolated states [14-13](#)

suspended states [14-13](#)

interface statistics

description [10-13](#)

interoperability

configuring interop mode 1 [22-8](#)

VSANs [15-12](#)

interop modes

configuring mode 1 [22-8](#)

default settings [22-12](#)

description [22-8](#)

IOD. See in-order delivery

ISLs

port channel links [14-1](#)

isolated VSANs

description [15-8](#)

displaying membership [15-9](#)

L

link costs

configuring for FSPF [18-6](#)

link failures

recovering [27-1](#)

load balancing

attributes [15-11](#)

attributes for VSANs [15-5](#)

configuring [15-11](#)

description [14-2](#), [15-11](#)

port channels [14-1](#)

logical unit numbers. See LUNs

M

MAC addresses

configuring secondary [22-5](#)

McData

native interop mode [22-8](#)

merged fabrics

autoreconfigured [11-7](#)

N

name servers

interoperability [22-9](#)

LUN information [21-1](#)

Send comments to nx5000-docfeedback@cisco.com

rejecting duplicate pWWNs [19-3](#)

NL ports

hard zoning [16-21](#)

NP links [12-2](#)

N ports

hard zoning [16-21](#)

zone enforcement [16-21](#)

zone membership [16-2](#)

See also Nx ports

NPV, configuring [12-4](#)

NPV mode [12-2](#)

NL ports [26-1](#)

Nx ports

FCS support [26-1](#)

See also N ports;

O

operational states

description [10-5](#)

P

passwords

DHCHAP [23-7, 23-8](#)

persistent FC IDs

configuring [11-16](#)

description [11-15](#)

enabling [11-16](#)

PLOGI

name server [19-3](#)

PortChannel Protocol

converting autocreated groups to manually configured [14-17](#)

port channel Protocol

autocreation [14-16](#)

creating channel group [14-15](#)

description [14-15](#)

port channel protocol

configuring autocreation [14-17](#)

enabling autocreation [14-17](#)

PortChannels

default settings [14-19](#)

verifying configurations [14-19 to ??](#)

port channels

adding interfaces [14-12, 14-13](#)

administratively down [10-7](#)

comparison with trunking [14-2](#)

compatibility checks [14-12](#)

compatibility with DHCHAP [23-3](#)

configuration guidelines [14-6](#)

configuring [14-12](#)

creating [14-10](#)

deleting [14-11](#)

deleting interfaces [14-14](#)

description [14-1](#)

forcing interface additions [14-13](#)

in-order guarantee [18-14](#)

interface states [14-13](#)

interoperability [22-9](#)

link changes [18-13](#)

link failures [18-2](#)

load balancing [14-2](#)

misconfiguration error detection [14-6](#)

port modes

auto [10-4](#)

port security

activating [24-7](#)

activation [24-2](#)

activation rejection [24-7](#)

adding authorized pairs [24-14](#)

auto-learning [24-2](#)

compatibility with DHCHAP [23-3](#)

configuration guidelines [24-3](#)

configuring CFS distribution [24-15 to 24-18](#)

configuring manually without auto-learning [24-12](#)

deactivating [24-7](#)

default settings [24-21](#)

Send comments to nx5000-docfeedback@cisco.com

- deleting entries from database (procedure) [24-15](#)
- disabling [24-5](#)
- displaying settings (procedure) [24-9](#)
- displaying statistics (procedure) [24-9](#)
- displaying violations (procedure) [24-10](#)
- enabling [24-5](#)
- enforcement mechanisms [24-2](#)
- fabric binding comparison [25-1](#)
- forcing activation [24-7](#)
- license requirement [24-1](#)
- preventing unauthorized accesses [24-1](#)
- WWN identification [24-13](#)
- port security auto-learning
 - authorization examples [24-12](#)
 - description [24-2](#)
 - device authorization [24-11](#)
 - disabling [24-11](#)
 - enabling [24-10](#)
 - guidelines for configuring with CFS [24-3](#)
 - guidelines for configuring without CFS [24-4](#)
- port security databases
 - copying active to config (procedure) [24-9](#)
 - interactions [24-18](#)
 - manual configuration guidelines [24-4](#)
 - merge guidelines [24-18](#)
 - reactivating [24-8](#)
 - scenarios [24-19](#)
- port tracking
 - default settings [27-7](#)
 - description [27-1](#)
 - enabling [27-3](#)
 - guidelines [27-2](#)
 - monitoring ports in a VSAN [27-6](#)
 - multiple ports [27-5](#)
- port world wide names. See pWWNs
- principal switches
 - assigning domain ID [11-9](#)
- pWWNs
 - configuring fcalias members [16-17](#)

- rejecting duplicates [19-3](#)
- zone membership [16-2](#)

R

RCFs

- description [11-3](#)
- incoming [11-6](#)
- rejecting incoming [11-6](#)

reason codes

- description [10-5](#)

reconfigure fabric frames. See RCFs

redundancy

- VSANs [15-4](#)

redundant physical links

- example (figure) [18-2](#)

Registered State Change Notifications. See RSCNs

related documents [1-iv](#)

retransmitting intervals

- configuring for FSPF [18-7](#)
- description [18-7](#)

route costs

- computing [18-6](#)

RSCNs

- default settings [19-8](#)
- displaying information [19-5](#)
- suppressing domain format SW-RSCNs [19-7](#)

RSCN timers

- configuring [19-8](#)

runtime checks

- static routes [18-11](#)

S

scalability

- VSANs [15-4](#)

SCSI LUNs

- customized discovery [21-2](#)

Send comments to nx5000-docfeedback@cisco.com

- discovering targets [21-1](#)
 - displaying information [21-3](#)
 - starting discoveries [21-1](#)
 - SD port mode
 - description [10-4](#)
 - interface modes [10-4](#)
 - secondary MAC addresses
 - configuring [22-5](#)
 - small computer system interface. See SCSI
 - soft zoning
 - description [16-21](#)
 - See also zoning
 - source IDs
 - exchange based [14-3](#)
 - flow based [14-3](#)
 - in-order delivery [18-12](#)
 - path selection [15-11](#)
 - SPAN destination port mode. See SD port mode
 - static routes
 - runtime checks [18-11](#)
 - storage devices
 - access control [16-1](#)
 - switch priorities
 - configuring [11-4](#)
 - default [11-4](#)
 - sWWNs
 - configuring for fabric binding [25-4](#)
-
- T
- TE port mode
 - classes of service [10-4](#)
 - description [10-4](#)
 - TE ports
 - fabric binding checking [25-2](#)
 - FCS support [26-1, 26-2](#)
 - FSPF topologies [18-1](#)
 - interoperability [22-9](#)
 - recovering from link isolations [16-23](#)
 - trunking restrictions [13-1](#)
 - timeout values. See TOVs
 - TOVs
 - configuring for a VSAN [22-3](#)
 - default settings [22-12](#)
 - interoperability [22-8](#)
 - ranges [22-1](#)
 - tracked ports
 - binding operationally [27-3](#)
 - traffic isolation
 - VSANs [15-4](#)
 - trunk-allowed VSAN lists
 - description [13-4 to 13-5](#)
 - trunking
 - comparison with port channels [14-2](#)
 - configuration guidelines [13-1](#)
 - description [13-1](#)
 - displaying information [13-6](#)
 - interoperability [22-8](#)
 - link state [13-3](#)
 - merging traffic [13-2](#)
 - restrictions [13-1](#)
 - trunking E port mode. See TE port mode
 - trunking ports
 - associated with VSANs [15-8](#)
 - trunking protocol
 - default settings [13-7](#)
 - default state [13-2](#)
 - description [13-2](#)
 - detecting port isolation [13-2](#)
 - trunk mode
 - configuring [13-3](#)
-
- U
- unique area FC IDs
 - configuring [11-17](#)
 - description [11-17](#)

Send comments to nx5000-docfeedback@cisco.com

V

VSAN IDs

- allowed list [13-7](#)
- description [15-5](#)
- multiplexing traffic [10-4](#)
- range [15-4](#)
- VSAN membership [15-4](#)

VSANs

- advantages [15-3](#)
- allowed-active [13-1](#)
- comparison with zones (table) [15-4](#)
- compatibility with DHCHAP [23-3](#)
- configuring [15-6](#)
- configuring allowed-active lists [13-6](#)
- configuring FSPF [18-3](#)
- configuring trunk-allowed lists [13-4 to 13-6](#)
- default settings [15-12](#)
- deleting [15-9](#)
- description [15-1 to 15-4](#)
- displaying configuration [15-12](#)
- domain ID automatic reconfiguration [11-7](#)
- FC IDs [15-1](#)
- FCS support [26-1](#)
- features [15-1](#)
- FSPF [18-4](#)
- FSPF connectivity [18-1](#)
- interop mode [22-9](#)
- isolated [15-8](#)
- load balancing attributes [15-5](#)
- mismatches [10-7](#)
- multiple zones [16-5](#)
- names [15-5](#)
- name server [19-2](#)
- port tracking [27-6](#)
- states [15-5](#)
- TE port mode [10-4](#)
- timer configuration [22-3](#)
- traffic isolation [15-3](#)

- trunk-allowed [13-1](#)
- trunking ports [15-8](#)

W

world wide names. See WWNs

WWNs

- displaying information [22-5](#)
- port security [24-13](#)
- secondary MAC addresses [22-5](#)
- suspended connections [10-7](#)

Z

zone aliases

- conversion to device aliases [17-7](#)
- importing [17-7](#)

zone attribute groups

- cloning [16-27](#)

zone databases

- migrating a non-MDS database [16-28](#)

zone members

- adding to zones [16-10](#)
- converting to pWWN members [16-20](#)
- displaying information [16-15](#)

zones

- access control [16-13](#)
- adding to zone sets [16-20](#)
- adding zone members [16-10](#)
- backing up (procedure) [16-25](#)
- changing from enhanced zones [16-30](#)
- cloning [16-27](#)
- comparison with device aliases (table) [17-2](#)
- comparison with VSANs (table) [15-4](#)
- configuring [16-20](#)
- configuring aliases [16-17](#)
- configuring fcaliases [16-17](#)
- default policies [16-2](#)

Send comments to nx5000-docfeedback@cisco.com

- displaying information [16-28 to ??](#)
- editing full zone databases [16-7](#)
- enforcing restrictions [16-20](#)
- exporting databases [16-23](#)
- features [16-1, 16-4](#)
- importing databases [16-23](#)
- membership using pWWNs [15-4](#)
- merge failures [10-7](#)
- renaming [16-26](#)
- restoring (procedure) [16-25](#)
- See also default zones
- See also enhanced zones
- See also hard zoning;soft zoning [16-21](#)
- See also zoning;zone sets [16-2](#)

zone server databases

- clearing [16-28](#)

zone sets

- activating [16-13](#)
- adding member zones [16-20](#)
- cloning [16-27](#)
- configuring [16-12 to 16-17](#)
- considerations [16-4](#)
- copying [16-24](#)
- creating [16-13, 16-20](#)
- displaying information [16-28 to ??](#)
- distributing configuration [16-21](#)
- enabling distribution [16-21](#)
- exporting [16-23](#)
- exporting databases [16-23](#)
- features [16-1](#)
- importing [16-23](#)
- importing databases [16-23](#)
- one-time distribution [16-22](#)
- renaming [16-26](#)
- See also active zone sets
- See also active zone sets;full zone sets [16-5](#)
- See also zones;zoning [16-2](#)

zoning

- description [16-1](#)