



Cisco DCNM SAN Client Online Help

Last Modified: 2020-12-22

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

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 NON-INFRINGEMENT 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.

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

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned 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. (1721R)

© 2020 Cisco Systems, Inc. All rights reserved.



CONTENTS

[Full Cisco Trademarks with Software License](#) ?

CHAPTER 1

[Device Manager Help](#) 1

[Physical](#) 1

[Inventory](#) 1

[Modules - Status and Config](#) 2

[Power Supplies](#) 2

[Temperature Sensors](#) 3

[Fan](#) 3

[Switches](#) 3

[ISLs](#) 5

[NP Link](#) 5

[ISL's Statistics](#) 5

[Hosts](#) 6

[Enclosures](#) 7

[Device Manager - Preferences](#) 7

[Interface](#) 8

[Virtual Interface Groups](#) 8

[Virtual FC Interfaces](#) 8

[Ethernet Interfaces](#) 9

[Virtual FC Ethernet](#) 10

[Quick Configuration Tool](#) 10

[Ethernet Interface](#) 11

[Ethernet Interfaces iSCSI](#) 12

[Ethernet Interfaces iSCSI TCP](#) 13

[Ethernet Interfaces VLAN](#) 14

Ethernet VLAN	14
FC Interface Monitor Traffic	15
FC Interface Monitor Protocol	15
FC Interface Monitor Discards	16
FC Interface Monitor Link Errors	16
FC Interface Monitor Frame Errors	16
FC Interface Monitor Class 2 Traffic	17
FC Interface Monitor Class 2 Errors	17
FC Interface Monitor FICON	18
Check Oversubscription	18
Virtual FC Interface Monitor Traffic	18
Virtual FC Interface Monitor Discards	18
Virtual FC Interface Monitor Errors	19
Ethernet Interface Dot3Stats	19
Interface Monitor	20
Ethernet PortChannels	20
Ethernet Interface Monitor iSCSI Connections	21
Ethernet Interface Monitor TCP	21
FCIP Monitor	21
Monitor SVC Interface	22
Monitor SVC NPorts	22
Monitor SVC Session FCP	23
Monitor SVC Session Other	23
FCIP Interfaces	24
System Timeout	24
Interface License	25
General	25
FC Interfaces General	26
FC Interfaces Rx BB Credit	30
FC Interfaces Other	31
FC Interfaces FLOGI	31
FC Interfaces ELP	32
FC Interfaces Trunk Config	34
FCIP Interfaces Trunk Failures	35

FC Interfaces IP	35
FC Interfaces Physical	35
FC Interfaces Capability	36
FC Interfaces FICON Peer	36
Interfaces NPorts (SVC)	37
Interfaces Sessions	37
IP Statistics TCP	37
Port Channels Ethernet Interfaces	38
Port Channels FC Interfaces	38
Port Channels General	39
FlexAttach Global	40
FlexAttach Virtual PWWN	40
FlexAttach Physical to Virtual WWNs	41
FIPS	41
FCIP FICON Configuration	41
Port Channels AutoCreate	42
SPAN Sessions	42
Span Global	42
SPAN Source Interfaces	42
Port Tracking Dependencies	42
Port Tracking Force Shut	43
Port Guard	43
Bandwidth Reservation: 48-Port 96-Gbps Fibre Channel module	43
Bandwidth Reservation: 48-Port 48-Gbps Fibre Channel module	43
Bandwidth Reservation: 24-Port 48-Gbps Fibre Channel module	44
Bandwidth Reservation: 48-Port 256-Gbps Fibre Channel module	44
Bandwidth Reservation: 32-Port 256-Gbps Fibre Channel module	44
DS-X9448-768K9 (Luke) Line Card Bandwidth Reservation	45
FC	45
VSAN General	45
VSAN Membership	46
VSAN Interop-4 WWN	47
VSAN Timers	47
VSAN Default Zone Policies	47

IVR Local Topology	47
IVR Fabric ID	48
IVR Default Fabric ID	48
IVR Action	48
IVR RDI VSANs	48
IVR Active Topology	49
IVR Zoneset Status	49
IVR Discrepancies	50
IVR Domains	50
IVR FCID	50
IVR Zoneset Active Zones	50
IVR Zoneset Active Zones Attributes	51
IVR Zoneset Name	51
DPVM Actions	51
DPVM Config Database	52
DPVM Active Database	52
Domain Manager Running	52
Domain Manager Configuration	53
Domain Manager Domains	54
Domain Manager Statistics	54
Domain Manager Interfaces	55
Domain Manager Persistent FcIds	55
Domain Manager Allowed DomainIds	55
Zoneset Active Zones	55
Zoneset Unzoned	56
Zoneset Status	56
Zoneset Policies	57
Zoneset Active Zones Attributes	57
Zoneset Enhanced	58
Zoneset Read Only Violations	58
Zoneset Statistics	58
Zoneset LUN Zoning Statistics	59
Zoneset Members	59
Fabric Config Server Discovery	60

Fabric Config Server Interconnect Elements	60
Fabric Config Server Platforms (Enclosures)	60
Fabric Config Server Fabric Ports	61
FC Routes	61
FDMI HBAs	62
FDMI Ports	62
FDMI Versions	62
Flow Statistics	63
FCC	63
Diagnostics	64
FSPF General	64
FSPF Interfaces	65
FSPF Interface Stats	66
SDV Virtual Devices	67
SDV Real Devices	67
LUN Discover	68
LUN Targets	68
LUNs	69
Device Alias	69
Device Alias Configuration	69
Device Alias Mode	70
Device Alias Discrepancies	70
Name Server General	70
Name Server Advanced	71
Name Server Proxy	71
Name Server Statistics	72
Preferred Path Maps and Routes	72
Preferred Path Maps Active	73
Preferred Path All Match Criteria	73
Preferred Path Active Match Criteria	73
Preferred Path All Sets	74
RSCN Nx Registrations	74
RSCN Multi-PID Support	74
RSCN Event	75

RSCN Statistics	75
Multicast Root	75
QoS Policy Maps	75
QoS Class Maps	76
QoS Match Statements	76
QoS Class Maps by Policy Maps	76
QoS Policy Maps by VSAN	76
QoS DWRR	77
QoS Rate Limit	77
Timers and Policies	77
WWN Manager	78
NPV Traffic Map	79
NPV Load Balance	79
NPV External Interface Usage	79
NP Link	79
FCoE	80
Config	80
VSAN-VLAN Mapping	80
VLAN-VSAN Mapping	80
FCoE Statistics	81
Ficon	82
FICON VSANs	82
FICON VSANs Files	83
Global	83
FICON Port Attributes	83
FICON Port Configuration	84
FICON Port Numbers	85
FICON VSANs Director History	85
Fabric Binding Actions	85
Fabric Binding Config Database	86
Fabric Binding Active Database	86
Fabric Binding Database Differences	86
Fabric Binding Violations	87
Fabric Binding Statistics	88

Fabric Binding EFMD Statistics	88
IP Storage	89
FCIP Profiles	89
FCIP Tunnels	89
FCIP Tunnels (Advanced)	90
FCIP Tunnels (FICON TA)	91
FCIP Tunnels Statistics	91
FCIP XRC Statistics	91
iSCSI Connection	92
iSCSI Initiators	92
iSCSI Session Initiators	93
Module Control	94
iSCSI Global	94
iSCSI Session Statistics	94
iSCSI Targets	95
iSCSI iSLB VRRP	96
iSCSI Initiator Access	96
Initiator Specific Target	96
iSCSI Initiator PWWN	97
iSCSI Sessions	97
iSCSI Sessions Detail	97
IP Services	98
IP Routes	98
IP Statistics ICMP	98
IP Statistics IP	99
IP Statistics SNMP	101
IP Statistics UDP	102
mgmt0 Statistics	102
TCP UDP TCP	103
TCP UDP UDP	103
VRRP General	103
VRRP IP Addresses	104
VRRP Statistics	104
CDP General	105

CDP Neighbors	105
iSNS Profiles	106
iSNS Servers	106
iSNS Entities	106
iSNS Cloud Discovery	106
iSNS Clouds	107
iSNS Cloud Interfaces	107
Monitor Dialog Controls	107
iSNS Details iSCSI Nodes	108
iSNS Details Portals	109
Security	109
Security Roles	109
Security Role Rules	110
Feature Group Manager	110
AAA LDAP Servers	110
AAA Server Groups	111
AAA Search Map	112
AAA Applications	112
AAA Defaults	113
AAA General	113
AAA Statistics	114
iSCSI User	117
Common Roles	117
SNMP Security Users	117
SNMP Security Communities	118
Security Users Global	118
FC-SP General/Password	119
FC-SP Interfaces	120
FC-SP Local Passwords	120
FC-SP Remote Passwords	121
FC-SP Statistics	121
FC-SP SA (Security Association)	121
FC-SP ESP Interfaces	121
PKI General	122

PKI RSA Key-Pair	122
PKI Trust Point	123
PKI Trust Point Actions	124
PKI LDAP	124
PKI Certificate Map	125
PKI Certificate Map - Application	125
PKI Trust Point Detail	125
IKE Global	126
IKE Pre-Shared AuthKey	127
IKE Policies	127
IKE Initiator Version	127
IKE Tunnels	128
IPSEC Global	128
IPSEC Transform Set	128
IPSEC CryptoMap Set Entry	129
IPSEC Interfaces	129
IPSEC Tunnels	130
IP ACL Profiles	130
IP ACL Interfaces	130
IP Filter Profiles	130
SSH/Telnet	133
Port Security Actions	133
Port Security Config Database	135
Port Security Active Database	135
Port Security Database Differences	136
Port Security Violations	136
Port Security Statistics	137
IPsec	137
Events	137
Call Home General	137
Call Home Destinations	137
Call Home Email Setup	138
Call Home Alerts	138
Call Home HTTP Proxy Server	139

Call Home SMTP Servers	139
Call Home User Defined Command	139
Delayed Traps	139
Call Home Profiles	139
Event Destinations Addresses	140
Event Destinations Security (Advanced)	140
Event Filters General	141
Event Filters Interfaces	142
Event Filters Control	142
Link Incident History	142
RMON Thresholds Controls	143
RMON Thresholds 64bit Alarms	143
RMON Thresholds 32bit Alarms	145
RMON Thresholds Events	145
RMON Thresholds Log	146
Admin	146
Copy Configuration	146
Flash Files	146
Compact Flash	147
License Features	147
License Manager Keys	147
License Manager Install	148
License Manager Usage	149
Port Licensing	150
Feature Set	150
Feature Control	150
NTP Servers	150
NTP General	151
Running Processes	151
Show Startup/Running Config	152
Show EPLD Version	152
Copy Flash Files	152
Show Tech Support	153
Show Image Version	153

Show Onboard Log	153
Summary View	154
RLIR ERL	155
Preferred Host	155
Preferred Path	156
Edit iSCSI Advertised Interfaces	156
DNS General	156
DNS Servers	156
Cisco Fabric Services (CFS) Features	157
Cisco Fabric Services (CFS) IP Multicast	159
Cisco Fabric Service (CFS) IP Static Peers	159
Cisco Fabric Services (CFS) Feature by Region	159
Cisco Fabric Services (CFS) All Region	159
Cisco Fabric Services (CFS) Owner	160
Cisco Fabric Services (CFS) Merge	160
Logs	160
SysLog (Since Reboot)	160
SysLog (Severe Events)	161
Accounting Log	161
Switch Logging	161
Syslog Severity Levels	162
Syslog Servers	162
End Devices - Hosts	162
Intelligent Features – Summary	163
Data Mobility Manager – Modules	163
Storage Media Encryption	164
Members	164
Interfaces	164
Hosts	164
SSM Features	165
Summary	165
FCWA	165
SSM	166
MSM	166

SANTap CVT	166
SANTap DVT	166
NASB	167
NASB Target	167
Virtual Initiator	168
DMM Rate	168
FCWA Config Status	168
Statistics Status	168
Statistics I/O Traffic	169
Statistics I/O Traffic Details	169
Statistics SCSI Commands	170
Statistics SCSI Errors	170
Statistics SCSI Sense Errors	171
Compact	171

CHAPTER 2
Configuring Cisco DCNM SAN Server 173

Configuring Cisco DCNM-SAN Server	173
Information About Cisco DCNM-SAN Server	173
DCNM-SAN Server Features	173
Licensing Requirements For Cisco DCNM-SAN Server	174
Installing and Configuring Cisco DCNM-SAN Server	174
Installing Cisco DCNM-SAN Server	174
Data Migration in Cisco DCNM-SAN Server	177
Verifying Performance Manager Collections	177
Managing a Cisco DCNM-SAN Server Fabric	177
Selecting a Fabric to Manage Continuously	177
Cisco DCNM-SAN Server Properties File	178
Modifying Cisco DCNM-SAN Server	179
Adding Cisco DCNM-SAN Server Users	180
Removing Cisco DCNM-SAN Server Users	180
Changing the Cisco DCNM-SAN Server Username and Password	180
Changing the DCNM-SAN Server Fabric Discovery Username and Password	181
Changing the Polling Period and Fabric Rediscovery Time	181
Changing the IP Address of the Cisco DCNM-SAN Server	181

Using Device Aliases or FC Aliases	182
Server Federation	182
Restrictions	182
Mapping Fabric ID to Server ID	183
Opening the Fabric on a Different Server	183
Viewing the Sessions in a Federation	184
Additional References	184

CHAPTER 3

Configuring Authentication in Cisco DCNM-SAN 185

Configuring Authentication in Cisco DCNM-SAN	185
Information About Cisco DCNM-SAN Authentication	185
Best Practices for Discovering a Fabric	186
Setting Up Discovery for a Fabric	187
Performance Manager Authentication	187
Cisco DCNM-SAN Web Client Authentication	188

CHAPTER 4

Configuring Cisco DCNM-SAN Client 191

Configuring Cisco DCNM-SAN Client	191
Information About DCNM-SAN Client	191
Cisco DCNM-SAN Advanced Mode	191
Cisco DCNM-SAN Client Quick Tour: Server Admin Perspective	192
Cisco DCNM-SAN Main Window	192
Menu Bar	193
Tool Bar	193
Logical Domains Pane	194
Physical Attributes Pane	194
Information Pane	194
Fabric Pane	195
Cisco DCNM-SAN Client Quick Tour: Admin Perspective	195
Menu Bar	197
Toolbar	201
Logical Domains Pane	202
Physical Attributes Pane	203
Information Pane	206

Fabric Pane	207
Status Bar	212
Launching Cisco DCNM-SAN Client	212
Launching Fabric Manager Client in Cisco SAN-OS Release 3.2(1) and Later	212
Launching Fabric Manager Client in Cisco SAN-OS Release 3.2(1) and Later	214
Launching Cisco DCNM-SAN Client Using Launch Pad	214
Setting Cisco DCNM-SAN Preferences	215
Network Fabric Discovery	217
Network LAN Discovery	217
Viewing Ethernet Switches	217
Removing a LAN	218
Modifying the Device Grouping	218
Using Alias Names as Enclosures	218
Using Alias Names as Descriptions	219
Controlling Administrator Access with Users and Roles	219
Using Cisco DCNM-SAN Wizards	220
Cisco DCNM-SAN Troubleshooting Tools	220
Integrating Cisco DCNM-SAN and Data Center Network Management Software	221
Launching a Switch from the Topology Map	221

CHAPTER 5
Configuring Device Manager 223

Device Manager	223
Information About Device Manager	223
Device Manager Features	224
Using Device Manager Interface	224
Menu Bar	225
Toolbar Icons	226
Dialog Boxes	228
Tabs	228
Legend	228
Supervisor and Switching Modules	229
Context Menus	229
Launching Device Manager	230
Setting Device Manager Preferences	230

CHAPTER 6	Configuring Performance Manager	233
	Configuring Performance Manager	233
	Information About Performance Manager	233
	Data Interpolation	234
	Data Collection	234
	Using Performance Thresholds	234
	Flow Statistics	235
	Flow Setup Wizards	236
	Creating a Flow Using Performance Manager Flow Wizard	236

CHAPTER 7	Configuring High Availability	237
	Configuring High Availability	237
	About High Availability	237
	Switchover Processes	237
	Synchronizing Supervisor Modules	238
	Manual Switchover Guidelines	238
	Manually Initiating a Switchover	238
	Copying Boot Variable Images to the Standby Supervisor Module	239
	Displaying HA Status Information	239

CHAPTER 8	Configuring Trunking	241
	Configuring Trunking	241
	Information About Trunking	241
	Trunking E Ports	241
	Trunking F Ports	242
	Key Concepts	243
	Trunking Protocols	243
	Trunk Modes	244
	Trunk-Allowed VSAN Lists and VF_IDs	245
	Guidelines and Limitations	247
	General Guidelines and Limitations	247
	Upgrade and Downgrade Limitations	248
	Difference Between TE Ports and TF-TNP Ports	248

Trunking Misconfiguration Examples	249
Default Settings	249
Configuring Trunking	250
Enabling the Cisco Trunking and Channeling Protocols	250
Enabling the F Port Trunking and Channeling Protocol	250
Configuring Trunk Mode	250

CHAPTER 9

Configuring PortChannels 253

Configuring PortChannels	253
Information About PortChannels	253
PortChannels Overview	253
E PortChannels	254
F and TF PortChannels	254
PortChanneling and Trunking	254
Load Balancing	255
PortChannel Modes	257
PortChannel Deletion	258
Interfaces in a PortChannel	258
PortChannel Protocols	260
Channel Group Creation	261
Autocreation	262
Manually Configured Channel Groups	263
Prerequisites for PortChannels	263
Guidelines and Limitations	264
General Guidelines and Limitations	264
Generation 1 PortChannel Limitations	264
F and TF PortChannel Limitations	264
Valid and Invalid PortChannel Examples	265
Default Settings	266
Configuring PortChannels	266
Configuring PortChannels Using the Wizard	266
Configuring the PortChannel Mode	267
Deleting PortChannels	268
Adding an Interface to a PortChannel	268

Forcing an Interface Addition	269
Deleting an Interface from a PortChannel	269
Converting to Manually Configured Channel Groups	269

CHAPTER 10

Configuring N Port Virtualization 271

Configuring N Port Virtualization	271
Information About N Port Virtualization	271
NPV Overview	271
N Port Identifier Virtualization	271
N Port Virtualization	272
NPV Mode	273
NP Ports	274
NP Links	274
Default Port Numbers	275
NPV CFS Distribution over IP	275
NPV Traffic Management	276
Multiple VSAN Support	276
Guidelines and Limitations	277
NPV Guidelines and Requirements	277
NPV Traffic Management Guidelines	277
DPVM Configuration Guidelines	278
NPV and Port Security Configuration Guidelines	278
Configuring N Port Virtualization	278
Configuring NPV	278
Using the NPV Setup Wizard	279
Configuring NPV Traffic Management	282
Displaying the External Interface Usage for Server Interfaces	283

CHAPTER 11

Configuring Interfaces 285

Configuring Interfaces	285
Information About Interfaces	285
Interface Description	285
Interface Modes	285
Interface States	289

Graceful Shutdown	292
Port Administrative Speeds	293
Frame Encapsulation	293
Beacon LEDs	294
Speed LEDs	294
Bit Error Thresholds	294
SFP Transmitter Types	295
TL Ports	296
TL Port ALPA Caches	296
Port Guard	296
Port Monitor	297
Port Monitor Port Guard	298
Port Group Monitor	298
Local Switching	298
Slow Drain Device Detection and Congestion Avoidance	298
Management Interfaces	299
VSAN Interfaces	299
Prerequisites for Interfaces	300
Guidelines and Limitations	300
Generation 1 Interface Configuration Guidelines	300
Private Loop Configuration Guidelines	301
VSAN Interface Configuration Guidelines	301
Default Settings	301
Configuring Interfaces	302
Common Interface Configuration	302
Setting the Interface Administrative State	302
Configuring Interface Modes	302
Configuring 10-Gbps FC Mode	303
Configuring Port Administrative Speeds	303
Configuring Port Speed Group	304
Specifying a Port Owner Using DCNM-SAN	304
Specifying a Port Owner Using Device Manager	305
Configuring Beacon Mode	305
Configuring TL Ports	305

Configuring Port Guard Using DCNM-SAN	306
Configuring Port Guard Using Device Manager	306
Configuring Management Interfaces	307
Creating VSAN Interfaces	307
Configuring Average Credit Nonavailable Duration Threshold and Action	308
Verifying Interfaces Configuration	308
Displaying the Owned Ports	308
Obtaining Interface Statistics	309
Displaying SFP Transmitter Types	309
Monitoring a Port Group	309

CHAPTER 12

Configuration of Fibre Channel Interfaces 311

Configuration of Fibre Channel Interfaces	311
Information About Fibre Channel Interfaces	311
Generations of Modules and Switches	311
Port Groups	313
Port Rate Modes	315
Port Speed	320
Dynamic Bandwidth Management	320
Out-of-Service Interfaces	323
Oversubscription Ratio Restrictions	323
Bandwidth Fairness	328
Guidelines and Limitations	329
Combining Generation 1, Generation 2, Generation 3, and Generation 4 Modules	329
Local Switching Limitations	330
Port Index Limitations	330
PortChannel Limitations	333
Default Settings	336
Configuring Fibre Channel Interfaces	337
Task Flow for Migrating Interfaces from Shared Mode to Dedicated Mode	337
Task Flow for Migrating Interfaces from Dedicated Mode to Shared Mode	338
Task Flow for Configuring 12-Port 4-Gbps Module Interfaces	339
Task Flow for Configuring 4-Port 10-Gbps Module Interfaces	339
Reserving Bandwidth Quickly for the 8-Gbps Module Interfaces	340

Configuring Port Speed	341
Configuring Rate Mode	341
Configuring Local Switching	342
Disabling Restrictions on Oversubscription Ratios Using DCNM-SAN	342
Disabling Restrictions on Oversubscription Ratios Using Device Manager	343
Enabling Restrictions on Oversubscription Ratios Using DCNM-SAN	343
Enabling Restrictions on Oversubscription Ratios Using Device Manager	343
Enabling Bandwidth Fairness Using DCNM-SAN	344
Enabling Bandwidth Fairness Using Device Manager	344
Disabling Bandwidth Fairness Using DCNM-SAN	344
Disabling Bandwidth Fairness Using Device Manager	345
Taking Interfaces Out of Service	345
Releasing Shared Resources in a Port Group	345
Verifying Fibre Channel Interfaces Configuration	346
Displaying Diagnostics for Multiple Ports	346

CHAPTER 13
Using the CFS Infrastructure 349

Monitoring Network Traffic Using SPAN	349
Information About SPAN	349
SPAN Sources	349
IPS Source Ports	350
Allowed Source Interface Types	350
VSAN as a Source	350
SPAN Sessions	351
Specifying Filters	351
SD Port Characteristics	351
Monitoring Traffic Using Fibre Channel Analyzers	352
Single SD Port to Monitor Traffic	352
SD Port Configuration	353
Mapping the FC Tunnel	353
Creating VSAN Interfaces	353
Remote SPAN	353
Advantages of Using RSPAN	354
FC and RSPAN Tunnels	354

ST Port Configuration	354
ST Port Characteristics	354
Creating Explicit Paths	355
Guidelines and Limitations	355
SPAN Configuration Guidelines	355
Guidelines to Configure VSANs as a Source	355
Guidelines to Specifying Filters	356
RSPAN Configuration Guidelines	356
Default SPAN and RSPAN Settings	357
Configuring SPAN	357
Configuring SD Ports for SPAN	357
Configuring SD Ports for SPAN using DM	357
Configuring SPAN max-queued-packets	358
Creating SPAN Sessions	358
Configuring SPAN for Generation 2 Fabric Switches	359
Editing SPAN Sources	359
Deleting SPAN Sessions	360
Encapsulating Frames	360
Configuring Fibre Channel Analyzers Using SPAN	360
Configuring RSPAN	360
Configuring the Source Switch	361
Enabling FC Tunnels	361
Configuring All Intermediate Switches	361
Configuring VSAN Interfaces	361
Enabling IP Routing	361
Configuring the Destination Switch	362
Configuring the SD Port	362
Monitoring RSPAN Traffic	362
Configuration Examples for RSPAN	362
Single Source with One RSPAN Tunnel	362
Single Source with Multiple RSPAN Tunnels	362
Multiple Sources with Multiple RSPAN Tunnels	363
Field Descriptions for SPAN	363
SPAN Sessions	363

Span Global	364
SPAN Source Interfaces	364

CHAPTER 14
Configuring SNMP 365

Configuring SNMP	365
Information About SNMP Security	365
SNMP Version 1 and Version 2c	365
SNMP Version 3	366
SNMPv3 CLI User Management and AAA Integration	366
Restricting Switch Access	366
Group-Based SNMP Access	366
Creating and Modifying Users	366
AES Encryption-Based Privacy	367
Enabling SNMP Notifications	367
LinkUp/LinkDown Notifications for Switches	367
Default Settings	368
Configuring SNMP	369
Assigning SNMP Switch Contact and Location Information	369
Configuring SNMP Users from the CLI	369
Enforcing SNMPv3 Message Encryption	369
Enforce the SNMPv3 message encryption globally	370
Assigning SNMPv3 Users to Multiple Roles	370
Adding or Deleting Communities	371
Deleting a Community String	371
Configuring SNMP Trap and Inform Notifications	372
Configuring SNMPv2c Notifications	372
Configuring SNMPv3 Notifications	372
Enabling SNMP Notifications	373
Configuring the Notification Target User	376
Configuring LinkUp/LinkDown Notifications for Switches	376
Configuring Up/Down SNMP Link-State Traps for Interfaces	376
Configuring Entity (FRU) Traps	377
Configuring Event Security	377
Viewing the SNMP Events Log	377

Field Descriptions for SNMP	378
IP Statistics SNMP	378
SNMP Security Users	379
SNMP Security Communities	380
Security Users Global	380

CHAPTER 15

Configuring Domain Parameters 383

Configuring Domain Parameters	383
Information About Fibre Channel Domains	383
Domain Restart	384
Domain Manager Fast Restart	385
Switch Priority	385
fcdomain Initiation	385
Incoming RCFs	385
Autoreconfiguring Merged Fabrics	386
Domain IDs	386
Locking the Fabric	389
Committing Changes	389
Clearing a Fabric Lock	389
FC IDs	389
Guidelines and Limitations	391
Default Settings	391
Configuring Fibre Channel Domains	392
Configuring Domain Manager Turbo Mode	392
Restarting a Domain	393
Restarting a Domain	393
Configuring Switch Priority	393
Enabling or Disabling fcdomains	394
Configuring Fabric Names	394
Rejecting Incoming RCFs	394
Enabling Autoreconfiguration	395
Configuring Domain IDs	395
Specifying Static or Preferred Domain IDs	395
Configuring Allowed Domain ID Lists	395

Enabling Allowed Domain ID Distribution	396
Enabling Contiguous Domain ID Assignments	397
Configuring FC IDs	398
Enabling the Persistent FC ID Feature	398
Configuring Persistent FC IDs	398
Configuring Unique Area FC IDs for an HBA	399
Purging Persistent FC IDs	399
Clearing a Fabric Lock	400
Displaying Pending Changes	400
Displaying Session Status	400
Monitoring FC Domain	401
Displaying fcdomain Statistics	401
Field Descriptions for FC Domain	401
IVR Domains	401

CHAPTER 16
Configuring and Managing Zones 403

Configuring and Managing Zones	403
Information About Zoning	403
Zone Implementation	404
About the Edit Local Full Zone Database Tool	405
About Zone Sets	406
About Zone Set Creation	407
About the Default Zone	407
About FC Alias Creation	408
Zone Enforcement	409
Zone Set Distribution	409
About Recovering from Link Isolation	409
Zone Set Duplication	410
About Backing Up and Restoring Zones	410
About Zone-Based Traffic Priority	410
About Broadcast Zoning	411
About LUN Zoning	411
About Read-Only Zones	412
About Enhanced Zoning	412

Merging the Database	413
Smart Zoning	414
Licensing Requirements for Zoning	415
Guidelines and Limitations	415
Zone Member Configuration Guidelines	415
Active and Full Zone Set Considerations	415
Read-Only Zone Configuration Guidelines	416
Default Settings	416
Configuring Zones	417
Configuring a Zone	417
Configuring a Zone Using the Zone Configuration Tool	417
Adding Zone Members	418
Filtering End Devices Based on Name, WWN, or FC ID	419
Adding Multiple End Devices to Multiple Zones	419
Using the Quick Config Wizard	419
Configuring Zone Sets	421
Activating a Zone Set	421
Deactivating a Zone Set	421
Displaying Zone Membership Information	422
Configuring the Default Zone Access Permission	422
Creating FC Aliases	422
Adding Members to Aliases	423
Converting Zone Members to pWWN-based Members	423
Creating Zone Sets and Adding Member Zones	424
Filtering Zones, Zone Sets, and Device Aliases Based on Name	425
Adding Multiple Zones to Multiple Zone Sets	425
Enabling Full Zone Set Distribution	426
Enabling a One-Time Distribution	426
Importing and Exporting Zone Sets	427
Copying Zone Sets	427
Backing Up Zones	428
Restoring Zones	428
Renaming Zones, Zone Sets, and Aliases	429
Cloning Zones, Zone Sets, FC Aliases, and Zone Attribute Groups	430

Migrating a Non-MDS Database	430
Clearing the Zone Server Database	430
Configuring Zone-Based Traffic Priority	431
Configuring Default Zone QoS Priority Attributes	431
Configuring the Default Zone Policy	432
Configuring Smart Zoning	432
Configuring Global Zone Policies	433
Configuring Broadcast Zoning	433
Configuring a LUN-Based Zone	433
Assigning LUNs to Storage Subsystems	434
Configuring Read-Only Zones	434
Changing from Basic Zoning to Enhanced Zoning	435
Changing from Enhanced Zoning to Basic Zoning	435
Enabling Enhanced Zoning	435
Modifying the Zone Database	436
Analyzing a Zone Merge	436
Preventing Zones From Flooding FC2 Buffers	436
Broadcasting a Zone	437
Configuring System Default Zoning Settings	437
Configuring Zone Generic Service Permission Settings	437
Compacting the Zone Database for Downgrading	437
Displaying Zone Information	438
Configuration Examples for Zoning	438
Field Descriptions for Zones	438
Zone Set Active Zones	439
Zone Set Unzoned	439
Zone Set Status	439
Zone Set Policies	440
Zone Set Active Zones Attributes	440
Zone Set Enhanced	441
Smart Zoning	441
Zone Set Read Only Violations	442
Zone Set Statistics	442
Zone Set LUN Zoning Statistics	442

Zone Set Members 443

CHAPTER 17

Configuring FCoE 445

Configuring FCoE 445

About FCoE 445

Guidelines and Limitations 445

Configuring FCoE 445

Enabling FCoE 445

Configuring FCoE Using DCNM for SAN 446

Configuring FCoE Using Device Manager 447

Field Descriptions for FCoE 448

Feature Set 448

Control 448

Config 448

VSAN-VLAN Mapping 448

Additional References 449

Feature History for FCoE 449

CHAPTER 18

Configuring Dense Wavelength Division Multiplexing 451

Configuring Dense Wavelength Division Multiplexing 451

Information About DWDM 451

Configuring X2 DWDM Transceiver Frequency using DCNM Manager 451

Configuring X2 DWDM Transceiver Frequency using DCNM-SAN 452

Monitoring DWDM Links 452

Field Descriptions for DPVM 453

DPVM Actions 453

DPVM Config Database 453

DPVM Active Database 453

Additional References 454

CHAPTER 19

Configuring and Managing VSANs 455

Configuring and Managing VSANs 455

Information About VSANs 455

VSANs Topologies 456

VSAN Advantages	458
VSANs Versus Zones	458
VSAN Configuration	459
About VSAN Creation	460
About Port VSAN Membership	460
About the Default VSAN	460
About the Isolated VSAN	460
Operational State of a VSAN	461
About Static VSAN Deletion	461
About Load Balancing	462
About Interop Mode	462
About FICON VSANs	462
Host Provisioning Wizard	462
Licensing Requirements for VSAN	462
Default Settings	462
Configuring VSANs	463
Multi-tenancy with MDS9000 and DCNM for SAN	463
Creating VSANs	463
Assigning Static Port VSAN Membership	464
Deleting Static VSANs	464
Configuring Load Balancing	465
Commissioning a Host	465
Decommissioning a Host	466
Displaying Isolated VSAN Membership	467
Field Descriptions for VSAN	467
VSAN General	467
VSAN Membership	468
VSAN Interop-4 WWN	468
VSAN Timers	469
VSAN Default Zone Policies	469

CHAPTER 20
Discovering SCSI Targets 471

Discovering SCSI Targets	471
Information About SCSI LUN Discovery	471

Licensing Requirements for SCSI	472
Discovering SCSI Targets	472
Starting SCSI LUN Discovery using Device Manager	472
Initiating Customized Discovery using Device Manager	472
Field Descriptions for SCSI Targets	473
iSCSI Connection	473
iSCSI Initiators	473
iSCSI Targets	474
iSCSI Session Initiators	475
iSCSI Global	475
iSCSI Session Statistics	475
iSCSI iSLB VRRP	476
iSCSI Initiator Access	476
iSCSI Initiator PWWN	476
iSCSI Sessions	476
iSCSI Sessions Detail	477

CHAPTER 21

Configuring SAN Device Virtualization 479

Configuring SAN Device Virtualization	479
Information About SDV	479
Key Concepts	482
Automatic Failover and Fallback	482
Resolving Fabric Merge Conflicts	483
Licensing Requirements for SAN Device Virtualization	483
Guidelines and Limitations	483
SDV Requirements and Guidelines	484
Guidelines for Downgrading SDV	484
Downgrading with Attributes Configured	484
Downgrading with Virtual Initiators Configured	485
Downgrading with SDV LUN Zoning Configured	485
Default Settings	485
Configuring SDV	485
Configuring a Virtual Device	486
Example - Configuring a Zone for a Virtual Device	486

Linking a Virtual Device with a Physical Device	487
Field Descriptions for SDV	488
SDV Virtual Devices	488
SDV Real Devices	488
Additional References	488

CHAPTER 22

Configuring Fibre Channel Write Acceleration 491

Configuring Fibre Channel Routing Services and Protocols	491
Information About FSPF	491
FSPF Global Configuration	492
About SPF Computational Hold Times	492
About Link State Record Defaults	492
About FSPF Link Cost	492
About Hello Time Intervals	493
About Dead Time Intervals	493
About Retransmitting Intervals	493
About Disabling FSPF for Specific Interfaces	493
FSPF Routes	493
About Fibre Channel Routes	494
About Broadcast and Multicast Routing	494
About Multicast Root Switch	494
In-Order Delivery	495
About Reordering Network Frames	495
About Reordering PortChannel Frames	495
About Enabling In-Order Delivery	496
About Flow Statistics	496
Licensing Requirements for FSPF	497
Default Settings	497
Configuring FSPF	498
Configuring FSPF on a VSAN	498
Resetting FSPF to the Default Configuration	498
Enabling or Disabling FSPF	498
Configuring FSPF Link Cost	499
Configuring Hello Time Intervals	499

Configuring Dead Time Intervals	499
Configuring Retransmitting Intervals	500
Disabling FSPF for Specific Interfaces	500
Configuring Fibre Channel Routes	501
Setting the Multicast Root Switch	501
Enabling In-Order Delivery for a VSAN	501
Configuring the Drop Latency Time	502
Verifying FSPF Configuration	502
Displaying the FSPF Database	503
Displaying FSPF Statistics	503
Configuration Examples for FSPF	503
Fault Tolerant Fabric	504
Redundant Links	504
Failover Scenarios for PortChannels and FSPF Links	504
Field Descriptions for FSPF	505
FSPF General	505
FSPF Interfaces	506
FSPF Interface Stats	507
FSPF LSDB Links	508
FSPF LSDB LSRs	508
FSPF Statistics	508
Additional References	509

CHAPTER 23
Managing FLOGI, Name Server, FDMI, and RSCN Databases 511

Managing FLOGI, Name Server, FDMI, and RSCN Databases	511
Information About FLOGI	511
Name Server Proxy	511
About Registering Name Server Proxies	511
About Rejecting Duplicate pWWN	512
About Name Server Database Entries	512
FDMI	512
RSCN	512
About the multi-pid Option	513
RSCN Timer Configuration Distribution Using CFS	513

RSCN Timer Configuration Distribution	514
Locking the Fabric	514
Default Settings	515
Registering Name Server Proxies	515
Registering Name Server Proxies	515
Configuring the multi-pid Option	515
Suppressing Domain Format SW-RSCNs	516
Configuring the RSCN Timer with CFS	516
Configuring the RSCN Timer	516
Committing the RSCN Timer Configuration Changes	517
Discarding the RSCN Timer Configuration Changes	517
Clearing a Locked Session	517
Displaying FLOGI Details	517
Viewing Name Server Database Entries	517
Displaying RSCN Information	518
Field Descriptions for Databases	518
FC Interfaces FLOGI	518
FDMI HBAs	519
FDMI Ports	519
FDMI Versions	519
RSCN Nx Registrations	520
RSCN Multi-PID Support	520
RSCN Event	520
RSCN Statistics	520
Name Server General	521
Name Server Advanced	521
Name Server Proxy	522
Name Server Statistics	522

CHAPTER 24
Configuring FICON 523

Configuring FICON	523
Information About FICON	523
FICON Requirements	524
Cisco MDS-Specific FICON Advantages	524

FICON Cascading	528
FICON VSAN Prerequisites	528
FICON Port Numbering	528
Default FICON Port Numbering Scheme	529
Port Addresses	532
Implemented and Unimplemented Port Addresses	532
About the Reserved FICON Port Numbering Scheme	532
Installed and Uninstalled Ports	533
About Port Numbers for FCIP and PortChannel	533
FC ID Allocation	534
About Enabling FICON on a VSAN	534
FICON Information Refresh	535
About FICON Device Allegiance	535
Automatically Saving the Running Configuration	535
Port Prohibiting	536
About RLIR	536
FICON Configuration Files	537
Port Swapping	538
FICON Tape Acceleration	538
CUP In-Band Management	540
Licensing Requirements for FICON	540
Guidelines and Limitations	541
FICON Port Numbering Guidelines	541
Port Swapping Guidelines	541
FICON Tape Acceleration Configuration Guidelines	541
Default Settings	542
Configuring FICON	542
Assigning FICON Port Numbers to Slots	542
Reserving FICON Port Numbers for FCIP and PortChannel Interfaces	543
Setting Up a Basic FICON Configuration	543
Enabling FICON on a VSAN	547
Manually Enabling FICON on a VSAN	547
Deleting FICON VSANs	548
Suspending a FICON VSAN	548

Configuring the code-page Option	549
Assigning FC ID Last Byte	549
Allowing the Host to Move the Switch Offline	550
Allowing the Host to Change FICON Port Parameters	550
Allowing the Host to Control the Timestamp	550
Configuring SNMP Control of FICON Parameters	551
Automatically Saving the Running Configuration	551
Configuring FICON Ports	552
Binding Port Numbers to PortChannels	552
Binding Port Numbers to FCIP Interfaces	552
Configuring Port Blocking	552
Configuring the Default State for Port Prohibiting	553
Configuring Port Prohibiting	553
Assigning a Port Address Name	553
Specifying an RLIR Preferred Host	554
Clearing RLIR Information	554
Applying the Saved Configuration Files to the Running Configuration	555
Editing FICON Configuration Files	555
Copying FICON Configuration Files	555
Swapping Ports	556
Configuring FICON Tape Acceleration	556
Configuring FICON Tape Read Acceleration	556
Configuring XRC Acceleration	557
Configure XRC Acceleration	557
Configure XRC acceleration on an FCIP Tunnel Interface Using Device Manager	557
Place the CUP in a Zone	557
Calculating FICON Flow Load Balance	558
Receiving FICON Alerts	558
Viewing ESCON Style Ports	559
Displaying RLIR Information	559
Displaying FICON Configuration Files	559
Displaying XRC Acceleration Statistics	560
Displaying XRC Acceleration Statistics	560
Displaying FICON Port Address Information	560

Displaying IPL File Information	561
Viewing the History Buffer	561
Field Descriptions for FICON	561
FICON VSANs	561
FICON VSANs Files	562
Global	563
FICON Port Attributes	563
FICON Port Configuration	563
FICON Port Numbers	564
FICON VSANs Director History	564

CHAPTER 25

Creating Dynamic VSANs 565

Creating Dynamic VSANs	565
Information About DPVM	565
About DPVM Configuration	566
About DPVM Databases	566
About Autolearned Entries	566
About DPVM Database Distribution	567
About Locking the Fabric	567
About Copying DPVM Databases	568
Licensing Requirements for VSANs	568
Guidelines and Limitations	568
Default Settings	568
Creating DPVM	569
Configuring DPVM with the DPVM Wizard	569
Configuring DPVM Config and Pending Databases	570
Activating DPVM Config Databases	570
Enabling Autolearning	571
Clearing a Single Autolearned Entry	571
Clearing All Autolearned Entries	571
Disabling DPVM Database Distribution	572
Locking the Fabric	572
Committing Changes	572
Discarding Changes	573

Clearing a Locked Session	573
Copying DPVM Databases	573
Comparing Database Differences	574
Viewing the Pending Database	574
Field Descriptions for DPVM	574
DPVM Actions	575
DPVM Config Database	575
DPVM Active Database	575
Additional References	576

CHAPTER 26
Distributing Device Alias Services 577

Distributing Device Alias Services	577
Information About Device Aliases	577
About Device Alias Modes	577
Changing Mode Settings	577
Device Alias Mode Distribution	578
Merging Device Alias	578
Resolving Merge and Device Alias Mode Mismatch	578
Device Alias Features	579
Device Alias Requirements	579
Zone Aliases Versus Device Aliases	579
Device Alias Databases	580
About Device Alias Distribution	580
About Creating a Device Alias	580
Fabric Lock Override	580
About Legacy Zone Alias Configuration Conversion	581
Guidelines and Limitations	581
Default Settings	581
Configuring Device Aliases	581
Creating Device Aliases	582
Distributing the Device Alias Database	582
Committing Changes	582
Discarding Changes	583
Using Device Aliases or FC Aliases	583

Populating Device Alias to Interface Description	584
Rename Device Alias	584
Field Descriptions for Device Aliases	584
Device Alias Configuration	585
Device Alias Mode	585
Device Alias Discrepancies	585

CHAPTER 27

Configuring Advanced Fabric Features 587

Configuring Advanced Fabric Features	587
Information About Common Information Model	587
SSL Certificate Requirements and Format	587
Fibre Channel Time-Out Values	588
About fctimer Distribution	588
Fabric Lock Override	588
World Wide Names	589
Link Initialization WWN Usage	589
FC ID Allocation for HBAs	589
Default Company ID List	590
Switch Interoperability	590
About Interop Mode	591
Guidelines and Limitations	592
Default Settings	593
Configuring Timer Across All VSANs	594
Task Flow for Configuring Time Across All VSANs	594
Configuring Timer Per-VSAN	594
Enabling fctimer Distribution	595
Configuring a Secondary MAC Address	595
Configuring Interop Mode 1	596
Verifying the Company ID Configuration	596
Verifying Interoperating Status	597

CHAPTER 28

Configuring Users and Common Role 599

Configuring Users and Common Role	599
Information About Role-Based Authorization	599

About Roles	599
Rules and Features for Each Role	600
Rule Changes Between SAN-OS Release 3.3(1c) and NX-OS Release 4.2(1a) Affect Role Behavior	600
About the VSAN Policy	601
Role Distributions	601
About Role Databases	601
Locking the Fabric	602
About Common Roles	602
Mapping of CLI Operations to SNMP	602
Creating Users Guidelines	603
Characteristics of Strong Passwords	604
About SSH	604
Boot Mode SSH	604
SSH Authentication Using Digital Certificates	604
Passwordless File copy and SSH	605
Guidelines and Limitations	605
Default Settings	605
Configuring Users and Common Role	606
Configuring Roles and Profiles	606
Deleting Common Roles	607
Modifying Rules	607
Modifying the VSAN Policy	608
Committing Role-Based Configuration Changes	608
Discarding Role-Based Configuration Changes	608
Enabling Role-Based Configuration Distribution	609
Clearing Sessions	609
Configuring Users	609
Deleting a User	611
Configuring SSH Services	611
Generating the SSH Server Key Pair	611
Overwriting a Generated Key Pair	612
Enabling SSH or Telnet Service	613
Changing Administrator Password Using DCNM-SAN	613

Recovering the Administrator Password	614
Displaying Role-Based Information	614
Displaying Roles When Distribution is Enabled	615
Displaying User Account Information	615
Field Descriptions for Users and Common Role	616
Common Roles	616
Feature History for Users and Common Role	616

CHAPTER 29

Configuring Security Features on External AAA Server 617

Configuring Security Features on an External AAA Server	617
Information About Switch Management Security	617
Security Options	617
SNMP Security Options	618
Switch AAA Functionalities	618
LDAP	624
LDAP Authentication and Authorization	624
About RADIUS Server Default Configuration	624
About the Default RADIUS Server Encryption Type and Preshared Key	625
About RADIUS Servers	625
About Validating a RADIUS Server	625
About Vendor-Specific Attributes	626
VSA Format	626
Specifying SNMPv3 on AAA Servers	626
One-Time Password Support	627
About TACACS+	627
About TACACS+ Server Default Configuration	627
About the Default TACACS+ Server Encryption Type and Preshared Key	627
About TACACS+ Servers	628
Password Aging Notification through TACACS+ Server	628
About Validating a TACACS+ Server	629
About Users Specifying a TACACS+ Server at Login	629
About Bypassing a Nonresponsive Server	629
AAA Server Distribution	630
Starting a Distribution Session on a Switch	630

CHAP Authentication	630
MSCHAP Authentication	630
About Enabling MSCHAP	631
Local AAA Services	631
Accounting Services	631
Defining Roles on the Cisco Secure ACS 5.x GUI	631
Defining Custom Attributes for Roles	632
Guidelines and Limitations	632
Remote Authentication Guidelines	632
Guidelines and Limitations for LDAP	633
Merge Guidelines for RADIUS and TACACS+ Configurations	633
Default Settings	633
Configuring the RADIUS, TACACS+, and LDAP Server	635
Authorizing and Authenticating the Switch	635
Configuring Fallback Mechanism for Authentication	636
Configuring AAA Server Monitoring Parameters Globally	636
Configuring the Default RADIUS Server Encryption Type and Preshared Key	636
Setting the Default RADIUS Server Timeout Interval and Retransmits	637
Configuring an LDAP Server	637
Creating LDAP Search Map	638
Configuring a RADIUS Server	639
Validating a RADIUS Server	640
Allowing Users to Specify a RADIUS Server at Login	640
Setting the Default TACACS+ Server Encryption Type and Preshared Key	641
Setting the Default TACACS+ Server Timeout Interval and Retransmits	641
Configuring a TACACS+ Server	642
Allowing Users to Specify a TACACS+ Server at Login	643
Clearing TACACS+ Server Statistics	643
Configuring Server Groups	643
Enabling Radius Server Distribution	644
Enabling TACACS+ Server Distribution	645
Committing the Distribution	645
Discarding the Distribution Session	645
Clearing Sessions	646

Enabling MSCHAP Authentication	646
Configuring Cisco Access Control Servers	647
Verifying RADIUS and TACACS+ Configuration	651
Displaying RADIUS Server Statistics	652
Displaying TACACS+ Server Statistics	653
Displaying the Pending Configuration to be Distributed	653
Configuration Examples for LDAP	653

CHAPTER 30

Configuring Certificate Authorities and Digital Certificates 655

Configuring Certificate Authorities and Digital Certificates	655
Information About Certificate Authorities and Digital Certificates	655
Purpose of CAs and Digital Certificates	655
Trust Model, Trust Points, and Identity CAs	656
RSA Key-Pairs and Identity Certificates	656
Multiple Trusted CA Support	657
PKI Enrollment Support	657
Manual Enrollment Using Cut-and-Paste Method	657
Multiple RSA Key-Pair and Identity CA Support	658
Peer Certificate Verification	658
CRL Downloading, Caching, and Checking Support	658
OCSP Support	659
Import and Export Support for Certificates and Associated Key-Pairs	659
Maximum Limits	659
Default Settings	659
Configuring CAs and Digital Certificates	660
Generating an RSA Key Pair	660
Creating a Trust Point CA Association	660
Copying Files to Bootflash	661
Authenticating the CA	661
Confirming CA Authentication	662
Generating Certificate Requests	663
Installing Identity Certificates	663
Saving Your Configuration	664
Ensuring Trust Point Configurations Persist Across Reboots	664

Monitoring and Maintaining CA and Certificates Configuration	664
Exporting and Importing Identity Information in PKCS12 Format	665
Configuring a CRL	666
Deleting Certificates from the CA Configuration	666
Deleting RSA Key Pairs from Your Switch	667
Configuration Examples	667
Downloading a CA Certificate	669
Requesting an Identity Certificate	669
Revoking a Certificate	670
Generating and Publishing the CRL	670
Downloading the CRL	670
Importing the CRL	671

CHAPTER 31
Configuring FC-SP and DHCHAP 673

Configuring FC-SP and DHCHAP	673
Information About Fabric Authentication	673
DHCHAP	674
DHCHAP Compatibility with Existing Cisco MDS Features	674
About Enabling DHCHAP	674
About DHCHAP Authentication Modes	674
About the DHCHAP Hash Algorithm	675
About the DHCHAP Group Settings	675
About the DHCHAP Password	675
About Password Configuration for Remote Devices	676
About the DHCHAP Timeout Value	676
Enabling FC-SP on ISLs	676
Default Settings	677
Configuring DHCHAP	677
Enabling DHCHAP	677
Configuring the DHCHAP Mode	678
Configuring the DHCHAP Hash Algorithm	678
Configuring the DHCHAP Group Settings	678
Configuring DHCHAP Passwords for the Local Switch	679
Configuring DHCHAP Passwords for Remote Devices	679

Configuring the DHCHAP Timeout Value	679
Configuring DHCHAP AAA Authentication	680

CHAPTER 32

Configuring Cisco TrustSec Fibre Channel Link Encryption	681
Configuring Cisco TrustSec Fibre Channel Link Encryption	681
Information About Cisco TrustSec FC Link Encryption	681
Supported Modules	681
Cisco TrustSec FC Link Encryption Terminology	682
Support for AES Encryption	682
Guidelines and Limitations	682
Configuring Cisco TrustSec Fibre Channel Link Encryption	682
Setting Up Security Association Parameters using DCNM-SAN	682
Setting Up Security Association Parameters using Device Manager	683
Setting Up Security Association Parameters using Device Manager	684
Configuring ESP Settings	684
Configuring ESP Modes	684
Configuring ESP Using ESP Wizard	686
Verifying Cisco TrustSec Fibre Channel Link Encryption Configuration	686
Displaying FC-SP Interface Statistics	687
Displaying FC-SP Interface Statistics Using Device Manager	687

CHAPTER 33

Configuring FIPS	689
Configuring FIPS	689
Information About FIPS Self-Tests	689
Guidelines and Limitations	690
Enabling FIPS Mode using DCNM-SAN	690
Enabling FIPS Mode using DCNM Manager	691
Field Descriptions for FIPS	692
FIPS	692

CHAPTER 34

Configuring IPv4 and IPv6 Access Control Lists	693
Configuring IPv4 and IPv6 Access Control Lists	693
Information About IPv4 and IPv6 Access Control Lists	693
About Filter Contents	694

Protocol Information	694
Address Information	694
Port Information	695
ICMP Information	696
ToS Information	696
Guidelines and Limitations	696
Configuring IPv4-ACLs or IPv6-ACLs	697
Creating IPv4-ACLs or IPv6-ACLs	697
Creating IPv4-ACLs or IPv6-ACLs	698
Deleting IP-ACLs	699
Reading the IP-ACL Log Dump	700
Applying an IP-ACL to an Interface	700
Applying an IP-ACL to mgmt0	701
Configuration Examples for IP-ACL	702
Field Descriptions for IPv4 and IPv6 Access Control Lists	703
IP ACL Profiles	703
IP ACL Interfaces	704
IP Filter Profiles	704

CHAPTER 35
Configuring IPsec Network Security 707

Configuring IPsec Network Security	707
Information About IPsec Network Security	707
About IKE	709
IPsec Compatibility	709
IPsec and IKE Terminology	710
Supported IPsec Transforms and Algorithms	711
Supported IKE Transforms and Algorithms	712
About IPsec Digital Certificate Support	712
About IKE Initialization	715
About the IKE Domain	715
About IKE Tunnels	715
About IKE Policy Negotiation	715
Optional IKE Parameter Configuration	716
About Crypto IPv4-ACLs	717

About Transform Sets in IPsec	719
About Crypto Map Entries	720
About SA Lifetime Negotiation	721
About the AutoPeer Option	721
About Perfect Forward Secrecy	722
About Crypto Map Set Interface Application	722
IPsec Maintenance	722
Global Lifetime Values	723
Prerequisites for IPsec	723
Guidelines and Limitations	724
Crypto IPv4-ACL Guidelines	724
Crypto Map Configuration Guidelines	725
Default Settings	726
Enabling IPsec Using FCIP Wizard	726
Verifying IPsec and IKE	728
Configuring IPsec and IKE Manually	728
Using IPsec	728
Configuring an IKE Policy	729
Configuring the Keepalive Time for a Peer	729
Configuring the Initiator Version	729
Clearing IKE Tunnels or Domains	730
Refreshing SAs	730
Configuring Crypto	730
Configuring Transform Sets	731
Creating Crypto Map Entries	731
Setting the SA Lifetime	732
Configuring Perfect Forward Secrecy	732
Applying a Crypto Map Set	732
Configuring Global Lifetime Values	733
Field Descriptions for IPsec	733
IPsec	733
IKE Global	733
IKE Pre-Shared AuthKey	733
IKE Policies	734

IKE Initiator Version	734
IKE Tunnels	734
IPSEC Global	735
IPSEC Transform Set	735
IPSEC CryptoMap Set Entry	735
IPSEC Interfaces	736
IPSEC Tunnels	736

CHAPTER 36
Configuring Port Security 737

Configuring Port Security	737
Information About Port Security	737
Port Security Enforcement	737
About Auto-Learning	738
Port Security Activation	738
Database Activation Rejection	739
About Enabling Auto-learning	739
Auto-learning Device Authorization	739
Authorization Scenarios	740
About WWN Identification	741
Activation and Auto-learning Configuration Distribution	741
Database Interaction	743
Guidelines and Limitations	743
Database Merge Guidelines	743
Default Settings	743
Configuring Port Security	744
Configuring Port Security with Auto-Learning and CFS Distribution	744
Configuring Port Security with Auto-Learning without CFS	745
Configuring Port Security with Manual Database Configuration	745
Configuring Port Security Using the Configuration Wizard	745
Enabling Port Security	747
Activating Port Security	747
Activating the Port Security Forcefully	748
Reactivating the Database	748
Copying an Active Database to the Config Database	749

Configuring Auto-learning	749
Enabling Auto-learning	749
Disabling Auto-learning	750
Configuring Port Security Manually	750
Task Flow for Configuring Port Security	750
Adding Authorized Port Pairs	751
Deleting Port Security Setting	751
Configuring Port Security Distribution	752
Enabling Distribution	752
Locking the Fabric	753
Committing the Changes	753
Discarding the Changes	753
Interacting with the Database	753
Copying the Port Security Database	753
Deleting the Port Security Database	754
Cleaning the Port Security Database	755
Cleaning the Port Security Database	755
Displaying Activated Port Security Settings	756
Displaying Port Security Statistics	756
Displaying Port Security Violations	757
Field Descriptions for Port Security	757
Port Security Actions	757
Port Security Config Database	758
Port Security Active Database	758
Port Security Database Differences	759
Port Security Violations	759
Port Security Statistics	760

CHAPTER 37
Configuring Fabric Binding 761

Configuring Fabric Binding	761
Information About Fabric Binding	761
Port Security Versus Fabric Binding	761
Fabric Binding Enforcement	762
Licensing Requirements for Fabric Binding	762

Default Settings	763
Configuring Fabric Binding	763
Enabling Fabric Binding	763

CHAPTER 38**Configuring FCIP 765**

Configuring FCIP	765
Information About FCIP	765
FCIP Concepts	766
FCIP High-Availability Solutions	768
Ethernet PortChannels and Fibre Channel PortChannels	771
FCIP Profile Configuration	771
Peers	771
Quality of Service	774
E Ports	774
FCIP Write Acceleration	775
FCIP Tape Acceleration	777
FCIP Compression	782
Default Settings	783
Configuring FCIP	784
Enabling FCIP	784
Modifying an FCIP Link	785
Creating FCIP Profiles	786
Checking Trunk Status	786
Launching Cisco Transport Controller	787
Configuring TCP Parameters	787
Assigning a Peer IP Address	790
Configuring Active Connections	791
Enabling Time Stamp Control	791
Configuring B Ports	791
Configuring FCIP Write Acceleration	792
Configuring FCIP Tape Acceleration	793
Field Descriptions for FCIP	793
FCIP Monitor	793
FCIP Interfaces	794

FCIP Interfaces Trunk Failures	794
FCIP FICON Configuration	795
FCIP Profiles	795
FCIP Tunnels	796
FCIP Tunnels (Advanced)	796
FCIP Tunnels (FICON TA)	797
FCIP Tunnels Statistics	797
FCIP XRC Statistics	798

CHAPTER 39

Configuring SAN Extension Tuner 799

Configuring the SAN Extension Tuner	799
Information About the SAN Extension Tuner	799
SAN Extension Tuner Setup	801
Data Pattern	801
Licensing Requirements for SAN Extension Tuner	801
Default Settings	802
Configuring the SAN Extension Tuner	802
Tuning the FCIP Link	802
Using the SAN Extension Tuner Wizard	803

CHAPTER 40

Configuring iSCSI 805

Configuring iSCSI	805
Information About iSCSI	805
About iSCSI Configuration Limits	807
Presenting Fibre Channel Targets as iSCSI Targets	808
Presenting iSCSI Hosts as Virtual Fibre Channel Hosts	810
Initiator Identification	810
Initiator Presentation Modes	810
Transparent Initiator Mode	811
WWN Assignment for iSCSI Initiators	812
Static Mapping	813
Proxy Initiator Mode	813
VSAN Membership for iSCSI	814
Advanced VSAN Membership for iSCSI Hosts	814

iSCSI Access Control	814
iSCSI Session Authentication	816
iSCSI Immediate Data and Unsolicited Data Features	816
About iSLB	819
About iSLB Initiators	820
iSLB Initiator Targets	820
iSLB Session Authentication	821
About Load Balancing Using VRRP	821
Changing iSCSI Interface Parameters and the Impact on Load Balancing	823
VRRP Load Balancing Algorithm For Selecting Gigabit Ethernet Interfaces	823
About iSLB Configuration Distribution Using CFS	823
Locking the Fabric	824
iSCSI High Availability	825
Multiple IPS Ports Connected to the Same IP Network	828
VRRP-Based High Availability	829
Ethernet PortChannel-Based High Availability	830
iSNS	831
About iSNS Client Functionality	831
About iSNS Server Functionality	832
iSNS Client Registration and Deregistration	832
Target Discovery	833
About Cloud Discovery	833
Licensing Requirements for iSCSI	834
Guidelines and Limitations	834
Default Settings	835
Configuring iSCSI	836
Enabling iSCSI	836
Creating iSCSI Interfaces	837
Using the iSCSI Wizard	837
Enabling Dynamic Mapping	838
Creating Static Mapping	838
Advertising Static iSCSI Targets	839
Specifying the Initiator Identification	840
Configuring the iSCSI Initiator Idle Timeout	840

Configuring Static Mapping	840
Making the Dynamic iSCSI Initiator WWN Mapping Static	841
Checking for WWN Conflicts	842
Configuring the Proxy Initiator	842
Configuring VSAN Membership for iSCSI Hosts	843
Configuring Default Port VSAN for iSCSI Interfaces	843
Adding iSCSI Initiator to the Zone Database	844
Configuring Access Control in iSCSI	844
Configuring AAA Authentication for an iSCSI User	845
Configuring Authentication Mechanism	845
Configuring iSLB	848
Configuring iSLB Using Device Manager	848
Configuring iSLB Initiator Names or IP Addresses	850
Making the Dynamic iSLB Initiator WWN Mapping Static	850
Assigning VSAN Membership for iSLB Initiators	850
Configuring and Activating Zones for iSLB Initiators and Initiator Targets	852
Restricting iSLB Initiator Authentication	852
Mutual CHAP Authentication	853
Configuring Load Balancing Using VRRP	853
Distributing the iSLB Configuration Using CFS	853
Enabling iSLB Configuration Distribution	853
Committing Changes to the Fabric	854
Discarding Pending Changes	854
Clearing a Fabric Lock	855
Creating a Static iSCSI Virtual Target	855
Enabling the Trespass Feature for a Static iSCSI	855
Configuring iSCSI Authentication	856
Configuring No Authentication	856
Configuring CHAP with Local Password Database	856
Configuring CHAP with External RADIUS Server	857
Creating an iSNS Client Profile	858
Deleting an iSNS profile	858
Tagging a profile to an interface	859
Untagging a profile from an interface	859

Configuring iSNS Servers	859
Configuring the ESI Retry Count	860
Configuring iSNS Cloud Discovery	861
Enabling iSNS Cloud Discovery	861
Initiating On-Demand iSNS Cloud Discovery	861
Configuring Automatic iSNS Cloud Discovery	862
Configuring iSNS Cloud Discovery Distribution	862
Configuring iSNS Cloud Discovery Message Types	862
Configuration Examples for iSCSI	862
Example 1	863
Example 2	863
Example 3	863
Example of VSAN Membership for iSCSI Devices	864
Example of an iSNS Server	865
iSCSI Transparent Mode Initiator Example	865
Target Storage Device Requiring LUN Mapping Example	869
Field Descriptions for iSCSI	872
Ethernet Interfaces iSCSI	873
Ethernet Interfaces iSCSI TCP	873
Ethernet Interface Monitor iSCSI Connections	874
iSCSI Connection	874
iSCSI Initiators	875
iSCSI Targets	876
iSCSI Session Initiators	876
iSCSI Global	877
iSCSI Session Statistics	877
iSCSI iSLB VRRP	877
iSCSI Initiator Access	878
iSCSI Initiator PWWN	878
iSCSI Sessions	878
iSCSI Sessions Detail	878
iSNS Details iSCSI Nodes	879
iSCSI User	879
Edit iSCSI Advertised Interfaces	879

CHAPTER 41**Configuring IP Services 881**

Configuring IP Services 881

Information About IP Services 881

Traffic Management Services 881

Management Interface Configuration 882

About the Default Gateway 883

IPv4 Default Network Configuration 883

IPFC 884

About IPv4 Static Routes 884

About Overlay VSANs 885

About VRRP 885

DNS Server Configuration 887

Guidelines and Limitations 887

Default Settings 887

Configuring IP Services 888

Configuring Management Interface 888

Configuring the Default Gateway 888

Configuring an IP route or identify the default gateway using Device Manager 889

Configuring Overlay VSANs 889

Configuring Multiple VSANs 891

Configuring VRRP 892

Adding and Deleting a Virtual Router 892

Virtual Router Initiation 893

Adding Virtual Router IP Addresses 893

Setting the Priority for the Virtual Router 893

Setting the Time Interval for Advertisement Packets 893

Configuring or Enabling Priority Preemption 893

Setting Virtual Router Authentication 894

Tracking the Interface Priority 894

Field Descriptions for IP Services 894

IP Routes 894

IP Statistics ICMP 895

IP Statistics IP 896

IP Statistics SNMP	897
IP Statistics UDP	898
mgmt0 Statistics	899
TCP UDP TCP	899
TCP UDP UDP	899
VRRP General	899
VRRP IP Addresses	900
VRRP Statistics	900
CDP General	901
CDP Neighbors	901
iSNS Profiles	902
iSNS Servers	902
iSNS Entities	903
iSNS Cloud Discovery	903
iSNS Clouds	903
iSNS Cloud Interfaces	903
Monitor Dialog Controls	904
iSNS Details iSCSI Nodes	904
iSNS Details Portals	905

CHAPTER 42
Configuring IP Storage 907

Configuring IP Storage	907
Information About IP Storage	907
IPS Module Upgrade	908
MPS-14/2 Module Upgrade	908
Supported Hardware	908
Gigabit Ethernet Interfaces for IPv4 Configuration	909
Basic Gigabit Ethernet Configuration	909
IPS Module Core Dumps	910
About VLANs for Gigabit Ethernet	911
Interface Subnet Requirements	911
Verifying Gigabit Ethernet Connectivity	912
Gigabit Ethernet High Availability	912
VRRP for iSCSI and FCIP Services	912

About Ethernet PortChannel Aggregation	913
CDP	914
Licensing Requirements for IP Storage	914
Guidelines and Limitations	914
Default Settings	915
Configuring IP Storage	915
Configuring IPS Core Dumps	915
Verifying IP Storage Configuration	916
Verifying Module Status	916
Field Descriptions for IP Storage	916
FCIP Profiles	917
FCIP Tunnels	917
FCIP Tunnels (Advanced)	918
FCIP Tunnels (FICON TA)	919
FCIP Tunnels Statistics	919
FCIP XRC Statistics	919
iSCSI Connection	920
iSCSI Initiators	920
iSCSI Targets	921
iSCSI Session Initiators	922
Module Control	922
iSCSI Global	922
iSCSI Session Statistics	923
iSCSI iSLB VRRP	923
iSCSI Initiator Access	923
Initiator Specific Target	923
iSCSI Initiator PWWN	924
iSCSI Sessions	924
iSCSI Sessions Detail	924
Additional References	925

CHAPTER 43
Configuring IPv4 for Gigabit Ethernet Interfaces 927

Configuring IPv4 for Gigabit Ethernet Interfaces	927
Information About IPv4	927

Interface Descriptions	928
Beacon Mode	928
About VLANs for Gigabit Ethernet	928
Interface Subnet Requirements	929
Licensing Requirements for IPv4 for Gigabit Ethernet Interfaces	929
Guidelines and Limitations	929
Default Settings	930
Configuring IPv4	930
Configuring Gigabit Ethernet Interface	930
Configuring Autonegotiation	931
Configuring the MTU Frame Size	931
Configuring Promiscuous Mode	932
Configuring the VLAN Subinterface	932
Additional References	932

CHAPTER 44

Configuring IPv6 for Gigabit Ethernet Interfaces	935
Configuring IPv6 for Gigabit Ethernet Interfaces	935
Information About IPV6	935
Extended IPv6 Address Space for Unique Addresses	936
IPv6 Address Formats	936
IPv6 Address Prefix Format	936
IPv6 Address Type: Unicast	936
IPv6 Address Type: Multicast	938
ICMP for IPv6	939
Path MTU Discovery for IPv6	940
IPv6 Neighbor Discovery	940
Router Discovery	942
IPv6 Stateless Autoconfiguration	942
Dual IPv4 and IPv6 Protocol Stacks	943
IPv6 Addressing and Enabling IPv6 Routing	944
Transitioning from IPv4 to IPv6	945
Guidelines and Limitations	945
Default Settings	946
Configuring Basic Connectivity for IPv6	946

Configuring IPv6 Addressing and Enabling IPv6 Routing	946
Configuring IPv6 Routing using Device Manager	947
Configuring IPv4 and IPv6 Protocol Addresses	947
Configuring Neighbor Discovery Parameters	947
Configuring a IPv6 Static Route	948

CHAPTER 45

Configuring SCSI Flow Services 949

Configuring SCSI Flow Services	949
Information About SCSI Flow Services	949
SCSI Flow Services Overview	949
SCSI Flow Specification Attributes	950
SCSI Flow Manager	950
SCSI Flow Configuration Client	951
SCSI Flow Data Path Support	951
Licensing Requirements for SCSI Flow Services	951
Guidelines and Limitations	951
Default Settings	952
Configuring SCSI Flow Services	952
Enabling SCSI Flow Services	952
Enabling Intelligent Storage Services	952
Configuring Fibre Channel Using DCNM-SAN	953
Disabling Intelligent Storage Services	953
Verifying SCSI Flow Services	954
Displaying SCSI Flow Services Information	954
File Description for SCSI Flow Services	956
SSM	956
Virtual Initiator	956

CHAPTER 46

Configuring SCSI Flow Statistics 957

Configuring SCSI Flow Statistics	957
Information About SCSI Flow Statistics	957
SCSI Flow Statistics Overview	957
SCSI Flow Specification Attributes	958
SCSI Flow Manager	958

SCSI Flow Configuration Client	959
SCSI Flow Data Path Support	959
Licensing Requirements for SCSI Flow Statistics	959
Default Settings	959
Configuring SCSI Flow Statistics	959
Enabling SCSI Flow Statistics	959
Clearing SCSI Flow Statistics	960
Field Descriptions for SCSI Flow Statistics	960
SSM	960
Virtual Initiator	961

CHAPTER 47
Configuring Fibre Channel Write Acceleration 963

Configuring Fibre Channel Write Acceleration	963
Information About Fibre Channel Write Acceleration	963
Licensing Requirements for Fibre Channel Write Acceleration	964
Default Settings	964
Configuring Fibre Channel Write Acceleration	964
Enabling Fibre Channel Write Acceleration	964
Filed Description for Fibre Channel Write Acceleration	965
FCWA	965
SSM	966
Virtual Initiator	966
FCWA Config Status	966

CHAPTER 48
Monitoring the Network 967

Monitoring the Network	967
Information About Network Monitoring	967
Monitoring Health and Events	967
Device Discovery	968
Topology Mapping	968
Inventory Management	970
Viewing Logs from Device Manager	970

CHAPTER 49
Monitoring Performance 971

Monitoring Performance	971
Information About Performance Monitoring	971
Real-Time Performance Monitoring	971
Historical Performance Monitoring	972
Configuring Performance Manager	972
Creating a Flow with Performance Manager	972
Creating a Collection with Performance Manager	972
Using Performance Thresholds	972
Configuring the Summary View in Device Manager	973
Configuring Per Port Monitoring using Device Manager	974
Displaying DCNM-SAN Real-Time ISL Statistics	975
Viewing Performance Statics Using DCNM-SAN	975
Displaying Performance Manager Reports	976
Displaying Performance Summary	976
Displaying Performance Tables and Details Graphs	976
Displaying Performance of Host-Optimized Port Groups	977
Displaying Performance Manager Events	977
Generating Performance Manager Reports	977
Generating Top10 Reports in Performance Manager	977
Generating Top10 Reports Using Scripts	977
Configuring Performance Manager for Use with Cisco Traffic Analyzer	978
Exporting Data Collections	980
Exporting Data Collections to XML Files	980
Exporting Data Collections in Readable Format	980
Exporting Data Collections in Readable Format	981
Analyzing SAN Health	981
Installing the SAN Health Advisor Tool	982

CHAPTER 50

Configuring Call Home	985
Configuring Call Home	985
Information About Call Home	985
Call Home Features	986
About Smart Call Home	986
Call Home Destination Profiles	988

Call Home Alert Groups	988
Call Home Message Level Feature	988
Syslog-Based Alerts	989
RMON-Based Alerts	989
General E-Mail Options Using HTTPS Support	989
Multiple SMTP Server Support	990
Periodic Inventory Notification	990
Duplicate Message Throttle	990
Call Home Configuration Distribution	990
Fabric Lock Override	991
Clearing Call Home Name Server Database	991
EMC E-mail Home Delayed Traps	991
Event Triggers	991
Call Home Message Levels	993
Message Contents	995
Guidelines and Limitations	1003
Call Home Database Merger Guidelines	1003
Default Settings	1004
Configuring Call Home	1004
Task Flow for Configuring Call Home	1005
Configuring Contact Information	1005
Enabling Call Home Function	1006
Configuring Destination Profiles	1006
Associating an Alert Group	1007
Customizing Alert Group Messages	1008
Configuring General E-Mail Options	1010
Configuring HTTPS Support	1010
Enable or Disable Transport Method	1010
Configuring an HTTP Proxy Server	1011
Configuring Call Home Wizard	1011
Task Flow for Configuring Call Home Wizard	1011
Launching Call Home Wizard	1012
Enabling Periodic Inventory Notifications	1012
Configuring Duplicate Message Throttle	1013

Enabling Call Home Fabric Distribution	1013
Call Home Communications Test	1014
Configuring Delayed Traps	1015
Enabling Delayed Traps Using Cisco Device Manager	1016
Viewing Event Filter Notification	1016
Field Descriptions for Call Home	1016
Call Home General	1016
Call Home Destinations	1017
Call Home SMTP Servers	1017
Call Home E-mail Setup	1017
Call Home Alerts	1018
Call Home User Defined Command	1018
Delayed Traps	1018
Call Home Profiles	1018
Event Destinations Addresses	1019
Event Destinations Security (Advanced)	1019
Event Filters General	1019
Event Filters Interfaces	1020
Event Filters Control	1021

CHAPTER 51

Configuring System Message Logging	1023
Configuring System Message Logging	1023
Information About System Message Logging	1023
Monitoring Syslog Server from DCNM-SAN	1026
System Message Logging	1026
SFP Diagnostics	1027
Outgoing System Message Logging Server Facilities	1027
System Message Logging Servers	1028
System Message Logging Configuration Distribution	1028
Fabric Lock Override	1029
Guidelines and Limitations	1029
Default Settings	1029
Configuring System Message Logging	1030
Task Flow for Configuring System Message Logging	1030

Enabling or Disabling Message Logging	1030
Configuring Console Severity Level	1031
Configuring Monitor Severity Level	1031
Configuring Module Logging	1031
Configuring Facility Severity Levels	1031
Sending Log Files	1032
Configuring System Message Logging Servers	1032
Configuring System Message Logging Servers	1033
Fabric Lock Override	1034
Verifying Syslog Servers from DCNM-SAN Web Server	1034
Monitoring Logs	1034
Viewing Logs from DCNM-SAN Web Server	1034
Viewing Logs from Device Manager	1034

CHAPTER 52
Scheduling Maintenance Jobs 1037

Scheduling Maintenance Jobs	1037
Information About the Command Scheduler	1037
Scheduler Terminology	1037
Licensing Requirements for Command Scheduler	1038
Guidelines and Limitations	1038
Default Settings	1038

CHAPTER 53
Configuring RMON 1039

Configuring RMON	1039
Information About RMON	1039
RMON Configuration Information	1039
RMON Configuration Using Threshold Manager	1040
RMON Alarm Configuration Information	1040
Default Settings	1041
Configuring RMON	1041
Configuring the RMON Traps in SNMP	1041
Enabling RMON Alarms by Port	1041
Enabling 32-Bit and 64-Bit Alarms	1042
Creating RMON Alarms	1042

Enabling 32-Bit RMON Alarms for VSANs	1043
Enabling 32-Bit and 64-Bit RMON Alarms for Physical Components	1044
Creating a New RMON from Device Manager Threshold Manager	1044
Managing RMON Events	1045
Managing RMON Alarms	1045
Viewing the RMON Log	1045
Field Descriptions for RMON	1046
RMON Thresholds Controls	1046
RMON Thresholds 64bit Alarms	1046
RMON Thresholds 32bit Alarms	1047
RMON Thresholds Events	1048
RMON Thresholds Log	1048

CHAPTER 54

Configuring Fabric Configuration Server 1049

Configuring Fabric Configuration Server	1049
Information About FCS	1049
Significance of FCS	1050
Default Settings	1051
Configuring FCS	1051
Specifying an FCS Name	1051
Creating an FCS Platform	1051
Displaying FCS Discovery	1052
Displaying FCS Interconnect Element	1052
Displaying FCS Fabric Ports	1052
Field Descriptions for FCS	1052

CHAPTER 55

Monitoring Network Traffic Using SPAN 1055

Monitoring Network Traffic Using SPAN	1055
Information About SPAN	1055
SPAN Sources	1055
IPS Source Ports	1056
Allowed Source Interface Types	1056
VSAN as a Source	1056
SPAN Sessions	1057

Specifying Filters	1057
SD Port Characteristics	1057
Monitoring Traffic Using Fibre Channel Analyzers	1058
Single SD Port to Monitor Traffic	1058
SD Port Configuration	1059
Mapping the FC Tunnel	1059
Creating VSAN Interfaces	1059
Remote SPAN	1059
Advantages of Using RSPAN	1060
FC and RSPAN Tunnels	1060
ST Port Configuration	1060
ST Port Characteristics	1060
Creating Explicit Paths	1061
Guidelines and Limitations	1061
SPAN Configuration Guidelines	1061
Guidelines to Configure VSANs as a Source	1061
Guidelines to Specifying Filters	1062
RSPAN Configuration Guidelines	1062
Default SPAN and RSPAN Settings	1063
Configuring SPAN	1063
Configuring SD Ports for SPAN	1063
Configuring SD Ports for SPAN using DM	1063
Configuring SPAN max-queued-packets	1064
Creating SPAN Sessions	1064
Configuring SPAN for Generation 2 Fabric Switches	1065
Editing SPAN Sources	1065
Deleting SPAN Sessions	1066
Encapsulating Frames	1066
Configuring Fibre Channel Analyzers Using SPAN	1066
Configuring RSPAN	1066
Configuring the Source Switch	1067
Enabling FC Tunnels	1067
Configuring All Intermediate Switches	1067
Configuring VSAN Interfaces	1067

Enabling IP Routing	1067
Configuring the Destination Switch	1068
Configuring the SD Port	1068
Monitoring RSPAN Traffic	1068
Configuration Examples for RSPAN	1068
Single Source with One RSPAN Tunnel	1068
Single Source with Multiple RSPAN Tunnels	1068
Multiple Sources with Multiple RSPAN Tunnels	1069
Field Descriptions for SPAN	1069
SPAN Sessions	1069
Span Global	1070
SPAN Source Interfaces	1070

CHAPTER 56
Monitoring System Processes and Logs 1071

Monitoring System Processes and Logs	1071
Information About System Processes and Logs	1071
Saving Cores	1071
Saving the Last Core to Bootflash	1071
First and Last Core	1072
Online System Health Management	1072
Loopback Test Configuration Frequency	1073
Loopback Test Configuration Frame Length	1073
Hardware Failure Action	1073
Performing Test Run Requirements	1073
Tests for a Specified Module	1073
Clearing Previous Error Reports	1074
Interpreting the Current Status	1074
On-Board Failure Logging	1075
Default Settings	1075
Core and Log Files	1076
Clearing the Core Directory	1076
Clearing the Core Directory	1076
Configuring System Health	1076
Task Flow for Configuring System Health	1076

Performing Internal Loopback Tests	1077
Performing External Loopback Tests	1077
Performing Serdes Loopbacks	1077
Configuring On-Board Failure Logging	1078
Verifying System Processes and Logs Configuration	1078
Displaying System Processes	1078
Displaying System Status	1079
Displaying Core Status	1079
Additional References	1080
Feature History for System Processes and Logs	1080

CHAPTER 57
Configuring QoS 1081

Configuring QoS	1081
Information About QoS	1081
Configuring QoS	1082
Information About Control Traffic	1082
Enabling or Disabling Control Traffic	1082
Information About Data Traffic	1082
Comparing VSAN Versus Zone-Based QoS	1084
Configuring Data Traffic	1084
Information About Class Map Creation	1085
Creating a Class Map	1085
Information About Service Policy Definition	1086
About Service Policy Enforcement	1086
About the DWRR Traffic Scheduler Queue	1086
Changing the Weight in a DWRR Queue	1087
Limiting Ingress Port Rate Limiting	1088
About Limiting Ingress Port Rate	1088

CHAPTER 58
Configuring Port Tracking 1089

Configuring Port Tracking	1089
Information About Port Tracking	1089
Guidelines and Limitations	1090
Default Settings	1090

Configuring Port Tracking	1091
Enabling Port Tracking	1091
Information About Configuring Linked Ports	1091
Information About Tracked Port	1092
Information About Tracking Multiple Ports	1093
Information About Monitoring Ports in a VSAN	1093
Information About Forceful Shutdown	1093
Forcefully Shutting Down a Tracked PortDetailed Steps	1094

CHAPTER 59

Configuring FlexAttach Virtual pWWN 1095

Configuring FlexAttach Virtual pWWN	1095
Information About FlexAttach Virtual pWWN	1095
FlexAttach Virtual pWWN	1095
Difference Between San Device Virtualization and FlexAttach Port Virtualization	1096
FlexAttach Virtual pWWN CFS Distribution	1096
Security Settings for FlexAttach Virtual pWWN	1096
Guidelines and Limitations	1096
Configuring FlexAttach Virtual pWWN	1097
Automatically Assigning FlexAttach Virtual pWWN	1097
Launching FlexAttach in DCNM-SAN	1097
Manually Assigning FlexAttach Virtual pWWN	1098
Mapping pWWN to Virtual pWWN	1098
Using the Server Admin FlexAttach Wizards	1099
Pre-Configuring FlexAttach for a New Server	1099
Moving a Server to Another Port or Switch	1101
Replacing a Server with Another Server	1102

CHAPTER 60

Configuring Interface Buffers 1107

Configuring Interface Buffers	1107
Information About Interface Buffers	1107
Buffer-to-Buffer Credits	1107
Performance Buffers	1108
Buffer Pools	1108
BB_Credit Buffers for Switching Modules	1109

BB_Credit Buffers for Fabric Switches	1118
Extended BB_Credits	1119
Buffer-to-Buffer Credit Recovery	1121
Buffer-to-Buffer State Change Number	1121
Receive Data Field Size	1122
Configuring Interface Buffers	1122
Configuring Buffer-to-Buffer Credits	1122
Configuring Performance Buffers	1123
Configuring Extended BB_credits	1123
Configuring Receive Data Field Size	1123

CHAPTER 61**Verifying Ethernet Interfaces 1125**

Verifying Ethernet Interfaces	1125
Information About Ethernet Interfaces	1125
Default Settings	1125
Verifying Ethernet Interfaces Configuration	1126
Displaying Interface Information Using DCNM-SAN	1126
Displaying Interface Information Using Device Manager	1126



CHAPTER 1

Device Manager Help

- [Physical](#), on page 1
- [Interface](#), on page 8
- [FC](#), on page 45
- [FCoE](#), on page 80
- [Ficon](#), on page 82
- [IP Storage](#), on page 89
- [IP Services](#), on page 98
- [Security](#), on page 109
- [Events](#), on page 137
- [Admin](#), on page 146
- [Logs](#), on page 160
- [End Devices - Hosts](#), on page 162
- [Intelligent Features – Summary](#), on page 163
- [Data Mobility Manager – Modules](#), on page 163
- [Storage Media Encryption](#), on page 164
- [SSM Features](#), on page 165

Physical

This section includes the physical attributes for the DCNM SAN setup:

Inventory

Field	Description
Name	Field Replaceable Unit (FRU) name.
ModelName	Model name identifier.
SerialNumber	Primary and secondary serial numbers.
HardwareRevision	Hardware revision.
SoftwareRevision	The release version of Cisco NX-OS software.
Alias	Alias name as specified by a network manager.

Field	Description
AssetID	User-assigned asset tracking identifier as specified by a network manager.

Modules - Status and Config

Field	Description
Name	Module description.
Module	Module name identifier.
OperStatus	Module's operational state.
Reset	Click to reboot the module.
RateModeOverSubscriptionLimit	Select this option to control the restriction on the oversubscription ratio on modules that support it. By default, the restriction is enabled. If you disable this option, all the interfaces on the module are capable of operating at maximum admin speed, regardless of the available bandwidth.
BandwidthFairnessConfig	Select this option to control bandwidth fairness on modules that support it. By default, bandwidth fairness is enabled.
BandwidthFairnessOper	Shows if bandwidth fairness is enabled or disabled. By default, bandwidth fairness is enabled.
X2 xcvrFrequency Config	Specifies the transceiver frequency of the module. <ul style="list-style-type: none"> • notApplicable - Select this when the module does not support this configuration. • xcvrFreqX2FC - Select this to set the module's FC transceiver frequency to 10 Gigabyte. • xcvrFreqX2Eth - Select this to set the module's Ethernet transceiver frequency to 10 Gigabyte.
ResetReasonDescription	Why module was last reset.
Local Switching Mode	Shows the status of the local switching modules.
StatusLastChangeTime	When OperStatus was changed.
Power Admin	Allows you to power on and off the Field Replaceable Unit (FRU).
Power Oper	Field Replaceable Unit (FRU) operational power state.
Current	Current supplied by the Field Replaceable Unit (FRU).

Power Supplies

Field	Description
Name	Power supply location.
TotalPowerAvailable	Shows the available power. In combined mode, the total available power is twice the lesser of the two power supplies.

Field	Description
Redundant/Combined	Select to determine how the power supplies are configured. Redundant mode provides a backup power supply if one should fail, but the total power available is less.
ModelName	The model identifier.
OperStatus	Operational power state.
TotalAvailable	Total power available for power supply usage. When Mode is redundant, the total power available will be the lesser power capacity of the power supplies. When Mode is combined, the total power available will be twice the lesser of the power capacities of the operating power supplies.
TotalReserved	Total current drawn by powered-on FRUs

**Note**

If the power supply to the Uros and Paradise is either interrupted or turned off, the OperStatus in the power supply table displays "offEnvOther". However, the corresponding entry for the powered down device the inventory table will remain.

Temperature Sensors

Field	Description
Name	Sensor location.
Threshold Major	Major temperature threshold.
Threshold Minor	Minor temperature threshold.
Current	Most recent measurement seen by the sensor.
Status	The present operational status of the sensor.

Fan

Field	Description
Name	Fan location.
ModelName	The model identifier.
OperStatus	The current operating status.

Switches

Field	Description
Description	A description of the switch and software.

Field	Description
UpTime	The time since the network management portion of the system was last re-initialized.
Name	An administratively-assigned name for this switch.
Location	The physical location of this switch (e.g., 'telephone closet, 3rd floor').
Contact	The contact person for this switch, together with information on how to contact this person.
SwitchWWN	The World-Wide Name of this switch.
ClockDateAndTime	The current local date and time for the system. Setting this is equivalent to setting an automated clock and calendar.
TimeZone	The current local time zone for the system. The time zone must be entered in the format GMT, which is the number of hours difference between your time zone and GMT (Greenwich Mean Time).
ProcessorRAM	Total number of bytes of RAM available on the Processor.
NVRAM	Total number of bytes of NVRAM in the entity.
NVRAMUsed	Number of bytes of NVRAM in use.
FIPSMODEActivation	Enable or disable FIPS mode on the device. FIPS 140-2 is a set of security requirements for cryptographic modules and it details the U.S. Government requirements for cryptographic modules. A module will comprise both hardware and software such as a data center switching or routing module. The module is said to be in FIP- enabled mode when a request is recieved to enable the FIPS mode and a set of self-tests are successfully run in response to the request. If the self-tests fail, then an appropriate error is returned.
CPUUtilization	The average utilization of CPU on the active supervisor.
MemoryUtilization	The average utilization of memory on the active supervisor.
FlashPartitionSize	Flash partition size.
FlashPartitionFreeSpace	Free space within a Flash partition.
Status	The overall status of the switch.
Vendor	Switch vendor's name, such as Cisco, McData, or Brocade.
Model	Switch model name, such as MDS 9134 or MDS 9124.
Release	Switch software version.
NumFCPorts	Number of physical FC ports in the switch.
WWN	MAC address for the Ethernet VDCs that are discovered.
VDCId	Unique IDs for the Ethernet VDCs that are discovered.
FCoE Enabled	If true, FCoE is enabled for the Ethernet VDCs that are discovered.

ISLs

Field	Description
From Switch	The source switch of the link.
From Interface	The port index of source E_port of the link.
To Switch	The switch on the other end of the link.
To Interface	The port index of destination E_port of the link.
Status	The operational status of the link.

NP Link

Field	Description
NPIV (Core)	The NPIV core switch.
F Port	The connected F Port on the NPIV core switch.
NPV	NPV Switch.
NP Port	The connected port on the NPV switch.
Status	The operational status of the link.

ISL's Statistics

Field	Description
Description	An alias name for the interface, as specified by a network manager. For Port Channel and FCIP, this field will always show members if they are available. For FCIP, this field will show compress if compressed.
VSAN(s)	VSAN membership.
Mode	Operating mode of the port> (See Legend).
Connected To	Attached port. This could be a host, storage, or switch port.
Speed	Maximum bandwidth in Gbps.
Rx	One of the following: Utilization % Number of Bytes Number of Frames Average Frame Size
Rx Comp	The IP Compression ratio for received packets on the FCIP device.

Tx	One of the following: Utilization % Number of Bytes Number of Frames Average Frame Size
Tx Comp	The IP Copression ratio for transmitted packets on the FCIP device
Errors	Total number of Rx and Tx errors on the interface. Types of Rx errors include CRC errors, fragmented framed, unsupported class frames, runt frames, jabber frames, and giant Frames. Types of Tx errors are generally CRC errors, but these are rare. If the Errors field is not empty, there are probably Rx errors. For a more detailed breakdown of the error count, check the Monitor dialog box for appropriate interface.
Discards	Total number of Rx and Tx discards on the interface. Rx frames discarded are generally due to protocol errors. On rare occasions, a frame is received without any hardware errors, but a filtering rule set for the MAC address discards the frame due to a mismatch. Discarded Tx frames can be timeout frame discards (port is offline or not up), or timeout frames that are not sent back to the supervisor (class F/2 frames). If the Discards field is not empty, it is probably due to timeout frames.
Log	If checked, writes the record into the message log on each poll interval.

Hosts

Field	Description
Enclosure Name	The name of the enclosure.
Alias	The device alias of this entry.
Port WWN	The assigned PWWN for this host.
FcId	The FC ID assigned for this host.
Link Status	The operational status of the link.
Serial Number	Serial number.
Model	Model name.
Firmware Ver	The version of the firmware executed by this HBA.
Driver Ver	The version of the firmware executed by this HBA.
Information	The information list corresponding to this HBA.
Switch Interface	Interface on the switch that is connected with the end device.

Enclosures

Field	Description
IP Address	The IP address of the enclosure.
Elem. Mgr Use HTTP	Use HTTP to launch the local enclosure.
Elem. Mgr URL/Path	Use a URL to launch the local enclosure
Device Type	If host, it is HBA. If storage, it is the SCSI target.
Vendor	If host, it is HBA. If storage, it is the SCSI target.
Model	If host, it is HBA. If storage, it is the SCSI target.
Firmware Ver	The version of the firmware executed by this HBA.
Driver Ver	The version level of the driver software controlling this HBA.
OS	The type and version of the operating system controlling this HBA
Other Info	The information list corresponding to this HBA.

Device Manager - Preferences

Field	Description
Retry Requests # time(s) after #sec timeout	The number of retries to be attempted after time out (seconds).
Enable Status Polling	Check to enables status polling in every (specified number of) seconds
Trace SNMP packets in Message Log	Check to enable tracing SNMP packets in the message log.
Register for Events after Open, listen on Port 1163	Check to automatically register for events.
Show WWN Vendor	Check to enable showing the WWN vendor name. <ul style="list-style-type: none"> • Replace - Replace the existing vendor name with the new one. • Prepend - Attach the new vendor name to the beginning of the current vendor name.
Show Timestamps as Date/Time	Check for displaying the time stamp in the Date/Time format.
Telnet Path	Path to the telnet client.
Use Secure Shell instead of Telnet	Check to use secure shell.
CLI Session Timeout	Time interval for the CLI session (in seconds). Enter '0' to disable CLI timeout.
Show Tooltips in Physical View	Check to show tooltips.

Field	Description
Label Physical View Ports with	<ul style="list-style-type: none"> • FICON - Displays FICON as label for the ports on the device view. • Interface - Displays Interface as label for the ports on the device view.
Export Table	<ul style="list-style-type: none"> • Tab-Delimited - Exports the table to tab-delimited text file. • XML - Exports the table to xml file.

Interface

The following sections:

Virtual Interface Groups

The Bound Ethernet Interface field in the table can be modified. The remaining fields are for information only.

Field	Description
Switch	Name of the switch hosting this virtual interface group (VIG).
VIG Id	Virtual interface group identifier.
Bound Eth Interface	Physical Ethernet interface associated with this VIG.
Virtual Eth Interfaces	The virtual Ethernet interface bound to this VIG.
Virtual FC Interfaces	The virtual FC interface bound to this VIG.
Operational Status	The current operational state of the VIG.
CreationTime	Date and time when the VIG was created.



Note This table applies only to N5k switches running version less than 4.0(1a).

Virtual FC Interfaces

The following fields in the table can be modified: Description, Bind Type, Bind Interface, Bind MAC Address, FCF Priority, VSAN ID Port, Mode Admin, Status Admin. The remaining fields are for information only.

Field	Description
Switch	Name of the switch hosting this interface.
Interface	Interface name.
Description	Text description of the interface as specified by a network manager.

Field	Description
VIG Id	Virtual interface group to which this virtual FC interface is bound.
Bind Type	The type of interface associated with this virtual FC interface - physical Ethernet interface or MAC address of the FCoE Node (ENode).
Bind Interface	Physical Ethernet interface or Ethernet port channel associated with this virtual FC interface.
Bind MACAddress	MAC address of an FCoE Node (ENode) or a remote Fibre Channel Forwarder (FCF) identified by the virtual FC interface.
FCF Priority	The FCoE Initialization Protocol (FIP) priority value advertised by the FCF to ENodes.
VSAN ID Port	VSAN ID to which this interface is statically assigned.
VSAN Id Dynamic	Index of the VSAN to which this interface is statically assigned.
Mode Admin	The port mode configured by the user. Virtual FC interfaces support only fabric port (F Port) mode.
Rate Mode	Specifies the interface as dedicated mode or shared mode.
Speed Oper	Operational speed.
Mode Oper	The current operating mode of the port.
Speed Admin	The port speed configured by the user.
Status Service	Specifies whether the interface is in service or out of service.
Status Admin	The desired state of the interface.
Status Oper	The current operational state of the interface.
Status FailureCause	The cause of current operational state of the port.
Status LastChange	When the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this value is N/A.



Note VIG Id field applies only to N5k switches running version less than 4.0(1a).

Ethernet Interfaces

The Description and Admin fields in the table can be modified. The remaining fields are for information only.

Field	Description
Switch	Name of the switch hosting this interface.
Interface	Interface name.
Description	Text description of the interface as specified by a network manager.
VIG Id	Virtual interface group to which this virtual interface is bound.

Field	Description
Bound Eth Interface	Physical Ethernet interface associated with this virtual Ethernet interface.
Admin	The desired state of the interface.
Oper	The current operational state of the interface.
LastChange	When the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this value is N/A.
CDP (Enable)	Indicates whether the Cisco Discovery Protocol is currently running on this interface.
Duplex Status	The current mode of operation of the MAC entity. The status 'unknown' indicates that the current duplex mode could not be determined.
Enable Link Trap	Specifies whether Link Up or Link Down traps should be generated for this interface.

**Note**

This table applies only to N5k switches running version less than 4.0(1a).

Virtual FC Ethernet

Field	Description
Switch	Name of the switch hosting this interface.
Interface	Displays the name of the vFC interface and its association with other interfaces.
Description	Text description of the interface as specified by a network manager.
Status Admin	The desired state of the interface.
Status Oper	The current operational state of the interface.
Speed Oper	Operational speed of the interface
LastChange	When the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this value is N/A.

Quick Configuration Tool

Field	Description
Show All Interfaces	Check this checkbox to show all the available interfaces including the interfaces that are not available for binding to a vFC.
Auto Assign vFC Id	Check this checkbox to select vFC Id automatically. If you do not select this option you must manually enter a valid vFC Id.

Field	Description
Switch Operational Type	Click Ethernet Switch if you are not configuring any Fibre Channel interfaces on the switch. Click FCoE Switch if you are configuring Fibre Channel and FCoE interfaces.
Interface	Name of the physical Ethernet interface. If you hover the cursor over a physical Ethernet interface, any associated virtual interfaces are displayed in the tooltip.
FCoE VLAN(VSAN)	FCoE VLAN (VSAN) mapping to be used by the interface.
Admin Mode	Admin mode of the vFC interface, i.e. F or E
Eth Only	Configures the physical Ethernet without any virtual interfaces. Click the Eth Only button in the column header to set all the interfaces to this value.
vEth Only	Configures the physical Ethernet to have an associated VIG and a virtual Ethernet interface. Click the vEth Only button in the column header to set all the interfaces to this value.
vFC Only	Configures the physical Ethernet to have an associated VIG and a virtual FC interface. Click the vFC Only button in the column header to set all the interfaces to this value.
vFC	Configures the physical Ethernet to have an associated VIG and a virtual FC interface. Click the vFC button in the column header to set all the interfaces to this value.
vEth + vFC	Configures the physical Ethernet to have an associated VIG, a virtual Ethernet interface and a virtual FC interface. Click the vEth + vFC button in the column header to set all the interfaces to this value.
Configure Action Status	Displays the current status of the requested configuration changes.



Note vEth only, vFC only, vEth + vFC columns are not applicable for N5K switches running version 4.0(1a)N1



Note vFC column is applicable only for N5K switches running version 4.0(1a)N1



Note For earlier configured ports, mapping details will not be displayed in VLAN(VSAN) Mapping column.

Ethernet Interface

Field	Description
Description	An 'alias' name for the interface as specified by a network manager.
Mtu	The size of the largest frame which can be sent/received on the interface, specified in bytes.

Field	Description
Speed Oper	Operational speed of the interface.
Speed Admin	<ul style="list-style-type: none"> • notApplicable - The Speed change is not applicable for that port. • oneGigSpeed - The IPStorage port is configured as 1G. • tenGigSpeed - The IPStorage port is configured as 10G.
Failure Cause	Causes of the failures.
PhysAddress	The interface's MAC address.
Status Admin	The desired state of the interface.
Status Oper	The current operational state of the interface.
LastChange	When the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this value is N/A.
ConnectorPresent	True if the connector is detected.
CDP (Enable)	An indication of whether the Cisco Discovery Protocol is currently running on this interface.
IscsiAuthMethod	The authentication method.
Promiscuous Mode	<p>Checking or unchecking this option dictates the destination of the packets/frames. If this option is checked, then this interface accepts packets/frames that are addressed to this station. If this option is not selected, then packets accepted by the station are transmitted on the media.</p> <p>Checking or unchecking this option does not affect the reception of broadcast and multicast packets/frames by the interface.</p>
AutoNegotiate	Select this option to enable auto negotiation.
Beacon Mode	In beacon mode, an interface LED is assigned a flashing mode for identification. Select this option to enable beacon mode.
IPAddress/Mask	IP address and subnet mask for the interface.



Note SAN Admin users cannot change the ethernet interfaces settings in Cisco Nexus 5000 Series switches using Device Manager.

Ethernet Interfaces iSCSI

Field	Description
Description	An 'alias' name for the interface as specified by a network manager.
Speed	Operational speed.
PhysAddress	The interface's WWN.

Field	Description
Admin	The desired state of the interface.
Oper	The current operational state of the interface.
LastChange	When the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this value contains a N/A value.
PortVSAN	The VSAN that the interface belongs to.
ForwardingMode	Use Store and Forward if the HBA has problems with Passthrough.
Initiator ID Mode	How the initiator is identified on this interface, either by its iSCSI name (name) or by its IP address (ipaddress).
Enable	The initiator proxy mode for this interface. If true, then all the initiators coming on this interface would use the initiator configuration provided by this interface. The initiator configuration include port WWN and node WWN.
Assignment	How the initiator proxy mode FC addresses are assigned. If 'auto', then the FC addresses are automatically assigned. If it is 'manual', then they have to be manually configured.
Port WWN	The Port FC address used by the initiators on this interface when the initiator proxy mode is on.
Node WWN	The Node FC address used by the initiators on this interface when the initiator proxy mode is on.

Ethernet Interfaces iSCSI TCP

Field	Description
Local Port	Local interface TCP port.
SACK	Indicates if the Selective Acknowledgement (SACK) option is enabled or not.
KeepAlive	The TCP keep alive timeout for this iSCSI interface. If the value is 0, the keep-alive timeout feature is disabled.
MinTimeout	The TCP minimum retransmit time.
Max	The TCP maximum retransmissions.
SendBufferSize	The TCP send buffer size.
MinBandwidth	The TCP minimum bandwidth.
MaxBandwidth	The TCP maximum bandwidth.
Estimated Round Trip	The estimated round trip delay of network pipe used for B-D product computation. The switch can use this to derive the TCP window to advertise.
QoS	The TCP QoS code point.
PMTU Enable	Indicates if the Path MTU discovery option is enabled or not.
PMTU Reset Timeout	The PMTU reset timeout.

Field	Description
Connections Normal	The number of normal iSCSI connections.
Connections Discovered	The number of discovery iSCSI connections.
CWM Enable	If true, congestion window monitoring is enabled. If false, it is disabled.
CWM Burst Size	The maximum burst sent after a TCP sender idle period.
Max Jitter	The maximum delay variation (not due to congestion) that can be experienced by TCP connections on this interface.
Port	The local TCP port of this interface.

Ethernet Interfaces VLAN

Field	Description
Switch	Name of the switch.
Interface	Name of the interface.
VLAN mode	The mode in which this VLAN is configured. Static—A port with static VLAN membership directly assigned to a single VLAN. Dynamic—A port with dynamic VLAN membership assigned to a single VLAN based on the content of packets received on the port via VQP queries to VMPS. multiVLAN—A port with multiple VLAN memberships that are directly assigned to one or more VLANs.
VLAN list	The list of VLANs which are allowed on the switch.

Ethernet VLAN

Field	Description
Switch	Name of the switch.
ID	Switch ID
Trunk Mode	Specifies whether the mode is access or trunk.
Trunk Status	Ttrunking status of the port.
Native VLAN	Native VLANs
Allowed VLAN List	The list of VLANs which are allowed to be received/transmitted on the port.
Active VLAN List	The list or range of VLANs that are active on the switch.

FC Interface Monitor Traffic

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
C3 Rx Bytes	The number of Class 3 bytes, including the frame delimiters, received by this port from its attached Nx_Port.
C3 Rx Frames	The number of Class 3 frames, including the frame delimiters, received by this port from its attached Nx_Port.
C3 Tx Bytes	The number of Class 3 bytes, including the frame delimiters, transmitted by this port to its attached Nx_Port.
C3 Tx Frames	The number of Class 3 frames, including the frame delimiters, transmitted by this port to its attached Nx_Port.
CF Rx Bytes	The number of Class F bytes, including the frame delimiters, received by this port from its attached Nx_Port.
CF Rx Frames	The number of Class F frames, including the frame delimiters, received by this port from its attached Nx_Port.
CF Tx Bytes	The number of Class F bytes, including the frame delimiters, transmitted by this port to its attached Nx_Port.
CF Tx Frames	The number of Class F frames, including the frame delimiters, transmitted by this port to its attached Nx_Port.

FC Interface Monitor Protocol

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
LRRIn	The number of Link reset responses received by the FC-port.
LRROut	The number of Link reset responses transmitted by the FC-port.
OlsIns	The number of Offline Sequence errors received by the FC-Port.
OlsOuts	The number of Offline Sequence errors issued by the FC-Port.
NOSIn	The number of Non-Operational Sequences received by the FC-port.
NOSOut	The number of Non-Operational Sequences transmitted by the FC-port.
LinkResetIns	The number of link reset protocol errors received by the FC-Port from the attached FC-port.
LinkResetOuts	The number of link reset protocol errors issued by the FC-Port to the attached FC-Port.
TxWaitCount	The number of times the FC-port waited due to lack of transmit credits.
RxBBCredit	The maximum number of receive buffers available for holding Class 2, Class 3 received from the logged-in Nx_Port. It is for buffer-to-buffer flow control in the incoming direction from the logged-in Nx_Port to FC-port.

Field	Description
TxBBCredit	The total number of buffers available for holding Class 2, Class 3 frames to be transmitted to the logged-in Nx_Port. It is for buffer-to-buffer flow control in the direction from FC-Port to Nx_Port.
BBCreditTransitionFromZero	The number of transitions of BB credit out of zero state.

FC Interface Monitor Discards

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
Class2	The number of Class 2 frames discarded by this port.
Class3	The number of Class 3 frames discarded by this port.
ClassF	The number of Class F frames discarded by this port.
EISL	The number of Enhanced Inter Switch Link (EISL) frames discarded by the FC-port. EISL frames carry an EISL header containing VSAN among other information.
InDiscards	The total number of inbound frames which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol.
OutDiscards	The total number of outbound frames which were chosen to be discarded even though no errors had been detected to prevent their being transmitted.

FC Interface Monitor Link Errors

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
LinkFailures	The number of link failures detected by the FC-Port.
SigLosses	The number of signal losses detected by the FC-Port.
SyncLosses	The number of loss of synchronization failures detected by the FC-Port.
InvalidTxWords	The number of invalid transmission words detected by the FC-Port.
DelimiterErrors	The number of Delimiter Errors detected by the FC-Port.
AddressIdErrors	The number of address identifier errors detected by the FC-Port.

FC Interface Monitor Frame Errors

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
InvalidCrcs	The number of invalid CRCs detected by the FC-Port. Loop ports should not count CRC errors passing through when monitoring.

Field	Description
ELPFailures	The number of Exchange Link Parameters Switch Fabric Internal Link service request failures detected by the FC-Port. This is applicable to only Interconnect_Port, which are E_Port or B_Port.
Frag	The number of fragmented frames received by the FC-port.
Runts	The number of frames received by the FC-port that are shorter than the minimum allowable frame length regardless if the CRC is good or not.
Jabbers	The number of frames received by the FC-port that are longer than a maximum frame length and also have a CRC error.
TooLong	The number of frames received by the FC-port where the frame length was greater than what was agreed to in FLOGI/PLOGI. This could be caused by losing the end of frame delimiter.
TooShort	The number of frames received by the FC-port where the frame length was less than the minimum indicated by the frame header (normally 24 bytes), but it could be more if the DFCTL field indicates an optional header should be present.
Unknown	The number of unknown class frames received by FC-port.
EOFa	The number of frames received by FC-port with EOF aborts.
Framing	The number of framing errors. This denotes that the FC-port detected an inconsistency of frame structure.

FC Interface Monitor Class 2 Traffic

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
In Octets/In Frames	The number of Class 2 frame bytes and frames, including the frame delimiters, received by this port from its attached Nx_Port.
Out Octets/Out Frames	The number of Class 2 frame bytes and frames, including the frame delimiters, delivered through this port to its attached Nx_Port.

FC Interface Monitor Class 2 Errors

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
BSY	The number of busy frame responses.
FRJT	The number of F_RJT frames generated by this port against Class 2 frames.
PBSY	The number of times that port busy was returned to this port as result of a class 2 frame that could not be delivered to the other end of the link. This occurs if the destination Nx_Port is temporarily busy. PBSY can only occur on SOFc1 frames (the frames that establish a connection).
PRJT	The number of times that port reject was returned to this port as a result of a class 2 frame that was rejected at the destination Nx_Port.

FC Interface Monitor FICON

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
FramePacingTime	Number of 2.5 microsecond units that frame transmission is blocked due to zero credit.
DispErrorsInFrame	Number of frames with disparity errors.
EOFErrs	Number of frames with EOF errors.
DispErrsOutOfFrame	Number of frames with OOF errors.
InvalidOrderSets	Number of invalid or unrecognizable Order Sets outside of frames.

Check Oversubscription

Field	Description
Interval	
Elapsed	Time elapsed.
Interface	Name of the interface
InOctectRate	
OutOctectRate	

Virtual FC Interface Monitor Traffic

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
RxBytes	The number of bytes, including the frame delimiters, received by this port from its attached N_Port.
RxFrames	The number of frames, including the frame delimiters, received by this port from its attached N_Port.
TxBytes	The number of bytes, including the frame delimiters, transmitted by this port to its attached N_Port.
TxFrames	The number of frames, including the frame delimiters, transmitted by this port to its attached N_Port.

Virtual FC Interface Monitor Discards

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
InDiscards	The total number of inbound frames which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol.

Field	Description
OutDiscards	The total number of outbound frames which were chosen to be discarded even though no errors had been detected to prevent their being transmitted.

Virtual FC Interface Monitor Errors

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
InErrors	The number of incoming errors detected by the virtual FC port.
OutErrors	The number of outgoing errors detected by the virtual FC port.

Ethernet Interface Dot3Stats

Field	Description
Interface	Name of the interface.
Alignment Errors	The count of frames received on a particular interface that are an integral number of octets in length but do not pass the FCS check.
FCS Errors	The count of frames received on a particular interface that are an integral number of octets in length but do not pass the FCS check. This count does not include frames received with frame-too-long or frame-too-short error.
Single Collision Frames	The count of successfully transmitted frames on a particular interface for which transmission is inhibited by a single collision.
Multiple Collision Frames	The count of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collisions.
SQE Test Errors	The number of times the PLS sublayer generated the SQE TEST ERROR message for a particular interface.
Deferred Transmissions	The count of the number of frames for which the first transmission attempt on a particular interface is delayed because the medium is busy.
Late Collisions	The number of times that a collision is detected on a particular interface later than one slot time into the transmission of a packet.
Excessive Collisions	The count of the number of frames for which transmission on a particular interface fails because of excessive collisions. This counter does not increment when the interface is operating in full-duplex mode.
Internal Mac Transmit Errors	The count of the number of frames for which transmission on a particular interface fails because of an internal MAC sublayer transmit error.
Carrier Sense Errors	The number of times that a carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular interface.
Frame Too Longs	The count of number of frames received on a particular interface that exceed the maximum permitted frame size.

Field	Description
Internal Mac Receive Errors	The count of number of frames for which reception on a particular interface fails because of an internal MAC sublayer receive error.
Symbol Errors	For an interface operating at 100 Mb/s, the number of times there was an invalid data symbol when a valid carrier was present

Interface Monitor

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
Rx Bytes	The total number of bytes received on the interface, including framing characters.
RxFrames	The number of frames received on the interface.
Rx Multicast Frames	(Nexus 5000 Series only) The number of multicast frames received on the interface.
Rx Broadcast Frames	(Nexus 5000 Series only) The number of broadcast frames received on the interface.
TxBytes	The total number of bytes transmitted out of the interface, including framing characters.
TxFrames	The total number of frames transmitted out of this interface.
Tx Multicast Frames	(Nexus 5000 Series only) The number of multicast frames transmitted out of this interface.
Tx Broadcast Frames	(Nexus 5000 Series only) The number of multicast frames transmitted out of this interface.
RxErrors	The number of inbound frames that contained errors preventing them from being deliverable to a higher-layer protocol.
TxErrors	The number of outbound frames that could not be transmitted because of errors.
RxDiscards	The number of inbound frames which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol.
TxDiscards	The number of outbound frames which were chosen to be discarded even though no errors had been detected to prevent their being transmitted.

Ethernet PortChannels

Field	Description
Description	Alias name for the interface as specified by a network manager.
Members	Members of this Ethernet port channel.
Oper Speed	Operational speed of the interface.
Mtu	The size of the largest frame which can be sent/received on the interface, specified in bytes.
PhysAddress	The interface's MAC address.

Field	Description
Status Admin	The desired state of the interface.
Status Oper	The current operational state of the interface.
LastChange	When the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this value is N/A.

Ethernet Interface Monitor iSCSI Connections

Field	Description
RxBytes	Total number of bytes received on an iSCSI session.
TxBytes	Total number of bytes transmitted on an iSCSI session.
IPSEC	A collection of objects for iSCSI connection statistics.

Ethernet Interface Monitor TCP

Field	Description
Opens	The number of times connections have been opened.
Accepts	The number of times connections have been accepted.
Failed	The number of times connections have failed.
RxResets	The number of times connections have been reset.
Est	The number of connections that have been established.
RxSegs	The total number of segments received on established connections, including those received in error.
TxSegs	The total number of segments sent, except for those containing retransmitted bytes.
ReTxSegs	The total number of segments retransmitted.
BadSegs	The total number of segments received in error (e.g., bad checksums).
TxSegResets	The number of segments sent containing the "reset" flag.
SplitSeg	The number of segments sent which were less than the minimum.
DupACKs	The number of duplicate ACKs received.
RxBytes	The number of header and data bytes received.
TxBytes	The number of header and data bytes sent.

FCIP Monitor

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
C3 Rx Bytes	The number of incoming bytes of data traffic.
C3 Tx Bytes	The number of outgoing bytes of data traffic.
CF Rx Bytes	The number of incoming bytes of control traffic.
CF Tx Bytes	The number of outgoing bytes of control traffic.
Rx Error	The number of inbound frames that contained errors preventing them from being deliverable to a higher-layer protocol.
Tx Error	The number of outbound frames that could not be transmitted because of errors.
RxDiscard	The number of inbound frames which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol.
TxDiscard	The number of outbound frames which were chosen to be discarded even though no errors had been detected to prevent their being transmitted.

Monitor SVC Interface

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
Rx Bytes	Number of incoming bytes.
Rx Frames	Number of incoming frames.
Tx Bytes	Number of outgoing bytes.
Tx Frames	Number of outgoing frames.
Rx Errors	Number of incoming errors.
Tx Errors	Number of outgoing errors.
Rx Discards	Number of incoming discards.
Tx Discards	Number of outgoing discards.

Monitor SVC NPorts

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
Rx Bytes	Number of incoming bytes on this virtual N-port.

Field	Description
Rx Frames	Number of incoming frames on this virtual N-port.
Tx Bytes	Number of outgoing bytes on this virtual N-port.
Tx Frames	Number of outgoing frames on this virtual N-port.
Rx Bytes	Number of incoming bytes on this virtual N-port.
Rx Frames	Number of incoming frames on this virtual N-port.
Tx Bytes	Number of outgoing bytes on this virtual N-port.
Tx Frames	Number of outgoing frames on this virtual N-port.

Monitor SVC Session FCP

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
Cmds	Number of incoming FCP Command frames in this session.
XferRdys	Number of incoming FCP Transfer Ready frames in this session.
DataFrames	Number of incoming FCP Data frames.
Status	Number of incoming FCP status frames.
DataBytes	Number of incoming FCP Data bytes.
OverRuns	Number of incoming FCP Overrun frames in this session.
UnderRuns	Number of incoming FCP Underrun frames in this session.
Cmds	Number of outgoing FCP Command frames in this session.
XferRdys	Number of outgoing FCP Transfer Ready frames in this session.
DataFrames	Number of outgoing FCP Data frames.
Status	Number of outgoing FCP status frames.
DataBytes	Number of outgoing FCP Data bytes.
OverRuns	Number of outgoing FCP OverRun frames in this session.
UnderRuns	Number of outgoing FCP UnderRun frames in this session.

Monitor SVC Session Other

The Monitor dialog boxes have special [Monitor Dialog Controls](#).

Field	Description
InELSFrames	Number of incoming Extended Link Service frames in this session.
InBLSFrames	Number of incoming Basic Link Service frames in this session.
OutELSFrames	Number of outgoing Extended Link Service frames in this session.
OutBLSFrames	Number of outgoing Basic Link Service frames in this session.
InAborts	Number of incoming aborted frames in this session.
OutAborts	Number of outgoing aborted frames in this session.
OpenXchanges	Number of Open Exchanges in this session.
InBadFc2Drops	Number of FC2 dropped frames in this session.
InBadFcPDrops	Number of FCP dropped frames.
InFCPDataExcess	Number of FCP Data Excess frames in this session.

FCIP Interfaces

Field	Description
Description	Alias name for the interface as specified by a network manager.
PortVsan	The VSAN ID to which this interface is statically assigned.
Oper Mode	The current operating mode of the port.
AutoChannelCreate	If checked, automatically create the PortChannel.
Admin	The desired state of the interface.
Oper Status	The current operational state of the interface.
FailureCause	The cause of current operational state of the port.
LastChange	When the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this value is N/A.
FICON Address	The FICON port address of this port.

System Timeout

If frames residing in the switch for a long time, they should be regarded as congestion drop. If there is continuously no tx/rx credits received, it should be regarded as no credit drop. You can configure the timeout value of congestion drop and no credit drop in the Device Manager client. To configure the slow port monitor timeout, please go to **Admin > System Timeout**.

Field	Description
E port Congestion Drop	Specify the time for E port congestion drop. Or click on default to input a default value. The unit is ms.

Field	Description
F port Congestion Drop	Specify the time for F port congestion drop. Or click on default to input a default value. The unit is ms.
F port NoCredit Drop	Specify the time for no credit drop. Click on disable if you do not want to drop the frames without tx/rx credits or click on default to input a default value. The unit is ms.
E Port slowport-monitor	Specify the slowport-monitor timeout value for E port. Click on disable to disable slowport monitoring. Or click on default to input a default value. The unit is ms.
F Port slowport-monitor	Specify the slowport-monitor timeout value for F port. Click on disable to disable slowport monitoring. Or click on default to input a default value. The unit is ms.



Note To configure the slow port monitor time out values from SAN client, go to **Physical Attributes > Switches > System> Timeout**.

Interface License

Field	Description
Type	Specifies the license that can be acquired for a given interface. Currently, the Port Activation license can be defined.
Config	Displays the license for which an interface is eligible. An interface which is not eligible for any type of license will not be displayed.
Oper	The current state of port license on the interface is displayed.

General

Field	Description
Description	An 'alias' name for the interface as specified by a network manager.
Mtu	The size of the largest frame which can be sent/received on the interface, specified in bytes.
Oper	Operational speed
PhysAddress	The interface's MAC address.

Field	Description
Admin	State of the admin.
Oper	The current operational state of the interface.
LastChange	When the interface entered its current operational state.
CDP	Enable or disable CDP.
Default Gateway	The IP address of the default gateway.

FC Interfaces General

The following variables are not supported by all interfaces: PortVSAN, Port Mode Admin and Oper, Admin Speed, and FailureCause.

Field	Description
Description	Alias name for the interface as specified by a network manager.
VSAN Id Port	VSAN ID to which this interface is statically assigned.
VSAN Id Dynamic	The VSAN ID that this interface has been dynamically assigned (see DPVM).
CDP (Enable)	An indication of whether the Cisco Discovery Protocol is currently running on this interface.
Promiscuous Mode	<p>Checking or unchecking this option dictates the destination of the packets/frames. If this option is checked, then this interface accepts packets/frames that are addressed to this station. If this option is not selected, then packets accepted by the station are transmitted on the media.</p> <p>Checking or unchecking this option does not affect the reception of broadcast and multicast packets/frames by the interface9.</p>
Auto Negotiate	An indication of whether auto-negotiation of speed and duplex mode should be used on this interface.
Beacon Mode	In beacon mode, an interface LED is assigned a flashing mode for identification. Select this option to enable beacon mode.

Field	Description
Mode Admin	

Field	Description
	<p>The port mode configured by the user. Modes are:</p> <ul style="list-style-type: none"> • auto - If the user configured the port as auto, then the port initialization scheme determines the mode of the port. • F Port - In fabric 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. • FL Port - In fabric loop port mode, an interface functions as a fabric loop port. This port may be connected to one or more NL ports (including FL ports in other switches) to form a public arbitrated loop. • E Port - In expansion 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. • FX Port - Interfaces configured as Fx ports can operate in either F port or FL port mode. The Fx port mode is determined during interface initialization depending on the attached N port or NL port. This administrative configuration disallows interfaces to operate in any other mode—for example, preventing an interface to connect to another switch. • SD Port - In SPAN destination port mode, an interface functions as a switched port analyzer (SPAN). The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic that passes through a Fibre Channel interface. • TL Port - In translatable loop port mode, an interface functions as a translatable loop port. It may be connected to one or more private loop devices. TL port mode is specific to Cisco MDS 9000 family switches and have similar properties as FL ports. • ST Port - In the SPAN tunnel port (ST port) mode, an interface functions as an entry point port in the source switch for the RSPAN Fibre Channel tunnel. The ST port mode and the remote SPAN (RSPAN) feature are specific to switches in the Cisco MDS 9000 Family. When configured in ST port mode, the interface cannot be attached to any device, and thus cannot be used for normal Fibre Channel traffic.

Field	Description
	<ul style="list-style-type: none"> • TE Port - In trunking E port mode, an interface functions as a trunking expansion port. It may be connected to another TE port to create an Extended ISL (EISL) between two switches. TE ports are specific to Cisco MDS 9000 family switches. • B Port - While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as Cisco's PA-FC-1G Fibre Channel port adapter, implement a bridge port model to connect geographically dispersed fabrics. The oper mode on this port type is "read only" and it cannot be set. • TF Port - Trunking f_Port • TNP Port - Trunking N Proxy port mode applicable only to N-port Virtualizer (NPV) • NP Port - N Proxy port mode applicable only to N-port Virtualizer (NPV)
Mode Oper	The current operating mode of the port.
SpeedGroup	<p>Specifies the current speed group configuration on the given interface.</p> <ul style="list-style-type: none"> • None—The interface speed group configuration on this interface is not applicable. It is a read-only value. • 10G—The interface speed group configuration on this interface is 10G. • 1/2/4/8G—The interface speed group configuration on this interface as 1G-2G-4G-8G.
Speed Admin	<p>The port speed configured by the user. The port speed values are auto, 1Gb, 2Gb, 4Gb, 8Gb, 10Gb, autoMax2G, and autoMax4G.</p> <p>Note On a Cisco Nexus 5000 Series switch that runs Cisco NX-OS Release 4.2(2), you can configure the 8-Gbps administrative speed only on an M1060 switch module. You can configure the speed to 1 Gbps, 2 Gbps, or 4 Gbps on all switch modules on a Cisco Nexus 5000 Series switch that runs Cisco NX-OS Release 4.2(2) or earlier releases.</p>
Speed Oper	Operational speed.
RateMode	Specifies the interface as dedicated mode or shared mode.
Status Service	Specifies whether the interface is in service or out of service.

Field	Description
Status Admin	The desired state of the interface.
Status Oper	The current operational state of the interface.
Status FailureCause	The cause of current operational state of the port.
StatusWasEnabled	If true, this port successfully completed a link initialization.
Status LastChange	When the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this value is N/A.
Port Owner	Administratively assigned name of the current owner of the interface resource.

FC Interfaces Rx BB Credit

Field	Description
Oper	The receive buffer-to-buffer credits configured for the operational port mode.
Model	The BB_Credit model used by the FC-port. The alternate BB_Credit management model can be used in the arbitrated loop topology to manage the flow of frames between the two ports participating in the current loop circuit. Since this is a characteristic of a physical port, this is not applicable for Port Channel ports.
Admin	The receive buffer-to-buffer credits configured for this port.
Extended	The extended BB credits that can be configured on an FC port (in the range 256 through 4095). The acceptable value depends on the BB credit configuration of other ports on the module. This value can only be modified on modules that support the extended BB credit feature.
AdminISL	The receive buffer-to-buffer credits configured for this port to be used if it is operating in xE_port mode.
AdminFx	The receive buffer-to-buffer credits configured for this port to be used if it is operating in Fx mode.
PerfBuffer Admin	The performance buffers configured for this port. These buffers in addition to the buffer-to-buffer credits are used to improve the performance of a port. If a value of 0 is set, then the module uses the built-in algorithm to calculate the number of performance buffers to be used.

Field	Description
PerfBuffer Oper	The performance buffers presently operational on this port.
Oper Rx	The maximum number of receive buffers available for holding Class 2, Class 3, Class F frames received from the peer Interconnect_Port.
Oper Tx	The total number of buffers available for holding Class 2, Class 3, Class F frames to be transmitted to the peer Interconnect_Port.
Current Rx	The current value of receive buffer-to-buffer credits for this port.
Current Tx	The current value of transmit buffer-to-buffer credits for this port.
BbScn Notify	Indicates whether the Buffer-to-buffer State Change Number (BB_SC_N) mode is enabled. If checked, BB_SC_N mode is enabled. If unchecked, BB_SC_N mode is disabled.

FC Interfaces Other

Field	Description
PortChannel Id	The port channel that this interface belongs to.
Fabric WWN	The world wide name given to this interface.
Mtu bytes	The size of the largest frame which can be sent/received on the interface, specified in bytes.
RxDataFieldSize bytes	The largest Data_Field size for an FT_1 frame that can be received by this port.
HoldTime us	The maximum time that the FC-Port shall hold a frame in the transmitter buffer before discarding it, if it is unable to deliver the frame.
Auto Port Channel	Check if you want the PortChannel to be created automatically.
FEC Admin	Specifies the port FEC state configured.
FEC Oper	Specifies the current operating FEC state of the port.

FC Interfaces FLOGI

Field	Description
FcId	The address identifier that has been assigned to the logged-in Nx_Port.
PortName	The world wide name of the logged-in Nx_Port.

Field	Description
NodeName	The world wide name of the Remote Node the logged-in Nx_Port belongs to.
Original PWWN	The original port WWN for this interface
Version	The version of FC-PH that the Fx_Port has agreed to support from the Fabric Login.
BBCredit Rx	The maximum number of receive buffers available for holding Class 2, Class 3 received from the logged-in Nx_Port. It is for buffer-to-buffer flow control in the incoming direction from the logged-in Nx_Port to FC-port.
BBCredit Tx	The total number of buffers available for holding Class 2, Class 3 frames to be transmitted to the logged-in Nx_Port. It is for buffer-to-buffer flow control in the direction from FC-Port to Nx_Port. The buffer-to-buffer flow control mechanism is indicated in the respective BbCreditModel.
CoS	The classes of services that the logged-in Nx_Port has requested the FC-Port to support and the FC-Port has granted the request.
Class2 RxDataSize	The Class 2 Receive Data Field Size of the logged-in Nx_Port. Specifies the largest Data Field Size for an FT_1 frame that can be received by the Nx_Port.
Class2 SeqDeliv	Whether the FC-Port has agreed to support Class 2 sequential delivery during the Fabric Login. This is meaningful only if Class 2 service has been agreed. This is applicable only to Fx_Ports.
Class3 RxDataSize	The Class3 Receive Data Field Size of the logged-in Nx_Port. Specifies the largest Data Field Size for an FT_1 frame that can be received by the Nx_Port.
Class3 SeqDeliv	Whether the FxPort has agreed to support Class 3 sequential delivery during the Fabric Login. This is meaningful only if Class 3 service has been agreed. This is applicable only to Fx_Ports.

FC Interfaces ELP

Field	Description
Neighbor Port	The port world wide name of the peer Interconnect_Port.
Neighbor Switch	The node world wide name of the peer Node.

Field	Description
BBCredit Rx	<p>The maximum number of receive buffers available for holding Class 2, Class 3, Class F frames received from the peer Interconnect_Port. It is for buffer-to-buffer flow control in the incoming direction from the peer Interconnect_Port to local Interconnect_Port.</p> <p>The buffer-to-buffer flow control mechanism is indicated in the respective BbCreditModel.</p>
BBCredit Tx	<p>The total number of buffers available for holding Class 2, Class 3, Class F frames to be transmitted to the peer Interconnect_Port. It is for buffer-to-buffer flow control in the direction from the local Interconnect_Port to peer Interconnect_Port.</p> <p>The buffer-to-buffer flow control mechanism is indicated in the corresponding BbCreditModel.</p>
CoS	The classes of services that the peer Interconnect_Port has requested the local Interconnect_Port to support and the local Interconnect_Port has granted the request.
Class2 SeqDeliv	Whether the local Interconnect_Port has agreed to support Class 2 sequential delivery during the Exchange Link Parameters Switch Fabric Internal Link Service request. This is meaningful only if Class 2 service has been agreed.
Class2 RxDataSize	The Class 2 Receive Data Field Size of the peer Interconnect_Port. Specifies the largest Data Field Size for an FT_1 frame that can be received by the Interconnect_Port. This is meaningful only if Class 2 service has been agreed.
Class3 SeqDeliv	Whether the local Interconnect_Port has agreed to support Class 3 sequential delivery during the Exchange Link Parameters Switch Fabric Internal Link Service request. This is meaningful only if Class 3 service has been agreed.
Class3 RxDataSize	The Class 3 Receive Data Field Size of the peer Interconnect_Port. Specifies the largest Data Field Size for an FT_1 frame that can be received by the Interconnect_Port. This is meaningful only if Class 3 service has been agreed.
ClassF X_ID	When true indicates that the peer Interconnect_Port supplying this parameter requires that an interlock be used during X_ID assignment in Class F. This is meaningful only if Class F service has been agreed.

Field	Description
ClassF RxDataSize	The Class F Receive Data Field Size of the peer Interconnect_Port. Class F service is always agreed between two Interconnect_Ports. Specifies the largest Data Field Size for an FT_1 frame that can be received by the Interconnect_Port.
ClassF ConSeq	The number of sequence status blocks provided by the Interconnect_Port supplying the parameters for tracking the progress of a sequence as a sequence recipient. The maximum number of concurrent sequences that can be specified is 255. A value of N/A in this field is reserved.
ClassF EECredit	The maximum number of Class F data frames which can be transmitted by an Interconnect_Port without receipt of accompanying ACK or Link_Response frames. The minimum value of end-to-end credit is one. The end-to-end credit field specified is associated with the number of buffers available for holding the Data_Field of a Class F frame and processing the contents of that Data_Field by the Interconnect_Port supplying the parameters.
ClassF OpenSeq	The open sequences per exchange shall specify the maximum number of sequences that can be open at one time at the recipient between a pair of Interconnect_Ports for one exchange. This value is used for exchange and sequence tracking.

FC Interfaces Trunk Config

Field	Description
Admin	<p>The trunking mode configured by the user.</p> <ul style="list-style-type: none"> When set to nonTrunk, the port negotiates and converts the link into non-trunking mode. This port and the peer port's OperTrunkMode will not carry multiple VSAN traffic. When set to trunk, the port negotiates and converts the link into trunking mode only if the peer port is trunk or auto. When set to auto, the port is willing to convert the link to a trunk link only if the peer port is trunk.
Oper	The current trunking mode of the port.

Field	Description
Allowed VSANs	The list of VSANs which are allowed to be received/transmitted on the port when the port is operating in trunking mode. Only ports operating in trunk mode can belong to multiple VSANs.
Up VSANs	The list of VSANs whose operational state is up, that this port is associated with. Only ports operating in trunk mode can be associated to multiple VSANs. This is applicable to only ports operating in trunk mode.

FCIP Interfaces Trunk Failures

Field	Description
FailureCause	An entry is shown in this table if there is an error in the trunk status for the given VSAN.

FC Interfaces IP

Field	Description
Switch	The name of the switch.
Ethernet Interface	A unique value that identifies the ethernet interface.
Ethernet Status	The current operational state of the ethernet interface.
Ethernet IP Address	The Internet address for this entity.
Peer IP Address	The Internet address for this entity
Port	The Port ID string as reported in the most recent CDP message.
Peer Interface	A unique value that identifies the peer interface on this device to which this link pertains.
Peer Device Id	The Peer Device ID string as reported in the most recent CDP message.
IP Security Enabled	Specifies whether the IP Security is turned on or not.

FC Interfaces Physical

Field	Description
BeaconMode	If enabled, an interface LED is put into flashing mode for easy identification of a particular interface.
ConnectorPresent	If true, there is a physical connector.
ConnectorType	The module type of the port connector.
TransmitterType	The technology of the port transceiver.

Field	Description
Vendor	The connector unit vendor.
PartNumber	The connector unit part number.
Revision	The port revision of the connector unit.
SerialNo	The serial number of the connector unit.

FC Interfaces Capability

Field	Description
FC-PH Vers Low	The lowest version of FC-PH that the FC-Port is capable of supporting.
FC-PH Vers High	The highest version of FC-PH that the FC-Port is capable of supporting.
RxDataSize Min	The minimum size in bytes of the Data Field in a frame that the FC-Port is capable of receiving from its attached FC-port.
RxDataSize Max	The maximum size in bytes of the Data Field in a frame that the FC-Port is capable of receiving from its attached FC-port.
HoldTime Min	The minimum holding time (in microseconds) that the FC-Port is capable of supporting.
HoldTime Max	The maximum holding time (in microseconds) that the FC-Port is capable of supporting.
CoS	The Bit mask indicating the set of Classes of Service that the FC-Port is capable of supporting.
ServiceStateCapable	Indicates whether this interface is capable of handling service state change.
PortRateMode Capable	Indicates whether this interface is capable of being configured as dedicated or shared port rate modes.
AdminRxBbCreditExtendedCapable	If true, it is capable of changing the extended buffer-to-buffer credits on the interface. The user can configure the object <code>fcIfAdminRxBbCreditExtended</code> on this interface
Class2Seq Deliv	The flag indicating whether or not the FC-Port is capable of supporting Class 2 Sequential Delivery.
Class3Seq Deliv	The flag indicating whether or not the FC-Port is capable of supporting Class 3 Sequential Delivery.

FC Interfaces FICON Peer

Field	Description
TypeNumber	The type number of the peer node. For example, the type number could be 002105.
SerialNumber	The sequence number assigned to the peer node during manufacturing. For example, the serial number could be 000000023053.

Field	Description
Tag	The identifier of the port in the peer node connected to this port.
FcId	Address Identifier assigned to NX-Port
Status	Specifies the status of the row, is valid, invalid or old.
Name	Name of this port.
Manufacturer	The name of the company that manufactured the peer node. For example, the manufacturer info could be HTC.
ModelNumber	The model number of the peer node. For example, the model number could be F20.
PlantOfMfg	The plant code that identifies the plant of manufacture of the peer node. For example, the plant code of manufacture could be 00.
UnitType	The type of the peer node that this port is communicating.
Alert	The type of link incident that occurred on this interface.

Interfaces NPorts (SVC)

Field	Description
Pwwn	The WWN (Worldwide Name) of the virtual N-port.
FcId	Fibre Channel Identifier of the virtual N-port.
State	The operational state of the virtual N-port.
DownReason	If the state of the N-port is 'down' as depicted by the instance of State, this value denotes the reason why this N-port is 'down'.

Interfaces Sessions

Field	Description
NportPwwn	The WWN of the N-port that belongs to this session.
PeerPwwn	The WWN of the remote N-port for this session.
PeerNwwn	The WWN of the remote N-port for this session.
PeerFcId	Fibre Channel Identifier of the remote port for this session.

IP Statistics TCP

Field	Description
AttemptFails	The number of times TCP connections have made a direct transition to the CLOSED state from either the SYN-SENT state or the SYN-RCVD state, plus the number of times TCP connections have made a direct transition to the LISTEN state from the SYN-RCVD state.
InErrs	The total number of segments received in error (e.g., bad TCP checksums).

ActiveOpens	The number of times TCP connections have made a direct transition to the SYN-SENT state from the CLOSED state.
PassiveOpens	The number of times TCP connections have made a direct transition to the SYN-RCVD state from the LISTEN state.
EstabResets	The number of times TCP connections have made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state.
InSegs	The total number of segments received, including those received in error. This count includes segments received on currently established connections.
OutSegs	The total number of segments sent, including those on current connections but excluding those containing only retransmitted bytes.
RetransSegs	The total number of segments retransmitted - that is, the number of TCP segments transmitted containing one or more previously transmitted bytes.
OutRsts	The number of TCP segments sent containing the RST flag.

Port Channels Ethernet Interfaces

Field	Description
Description	Alias name for the interface as specified by a network manager.
Mtu	The size of the largest frame which can be sent/received on the interface, specified in bytes.
PhysAddress	The interface's address at its protocol.
Admin	The desired state of the interface.
Oper	The current operational state of the interface.
LastChange	When the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this value is N/A.
IPAddress/Mask	The IP address and mask of the interface.
iSCSI AuthMethod	The authentication method for this interface.
iSNS ProfileName	The iSNS server profile name for this interface.

Port Channels FC Interfaces

Field	Description
PortVsan	VSAN to which this interface is statically assigned.
Description	Alias name for the interface as specified by a network manager.
Admin Mode	The port mode configured by the user. If the user configured the port as auto(1), then the port initialization scheme determines the mode of the port. In this case the user can look at OperMode to determine the current operating mode of port. Only auto(1) or ePort(4) is allowed.

Field	Description
Oper Mode	The current operating mode of the port.
Admin Speed	The port speed configured by the user.
Oper Speed	The interface's current bandwidth per second.
Admin Status	The desired state of the interface.
Oper Status	The current operational state of the interface.
FailureCause	The cause of current operational state of the port.
LastChange	The time the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the switch, then this is a zero or N/A value.

Port Channels General

Field	Description
Admin Mode	The channel mode desired by the network manager.
Oper Mode	The current operating channel mode of the port.
Force	<p>The method to add port(s) to a Port Channel port.</p> <ul style="list-style-type: none"> • If unchecked, then a compatibility check is done on the parameters of the port(s) being added to this Port Channel. The port(s) being added must have the same physical and configured parameters as the Port Channel port. • If checked, a compatibility check is done on only physical parameters. The port(s) being added to this Port Channel port must have same physical parameters. The operation will fail only if the physical parameters are not same. The configured parameters of the port(s) being added are overwritten by configured parameters of this Port Channel port.
MemberList By Interface	The list of the E_ports that are members of this Port Channel port.
MemberList By FICON	The list of the E_ports that are members of this Port Channel port.
MemberList LoadBalanced	Those ports which are actively participating in the PortChannel.

Field	Description
LastAction Status	The status of the last operation (add or remove a member) done to change the member list of a Port Channel Port. When no ports are added or the last operation is successful then this value is successful. If this value is failed then the user can look at LastAddStatusCause to find the reason of failure.
LastAction FailureCause	The cause of failure to last operation (add or remove a member) done to change the member list of a Port Channel port.
LastAction Time	The timestamp indicating the time of last action performed on this entry.
CreationTime	The timestamp of this entry's creation time.
FICON Address	The FICON port address. If empty, then this channel is not used by FICON. (This column is displayed if FICON is enabled. This column is grayed out if the Port Channel is auto-created.)

FlexAttach Global

Field	Description
VirtualPwwnauto	Enables automatic generation of Virtual WWNs on all the F_port interfaces. If the value of VirtualPWWNauto is set to 'true', the value of VirtualWWN Auto of all the entries in the VirtualWWN table is implicitly set to true.

FlexAttach Virtual PWWN

Field	Description
virtual pWWN	This is the virtual port WWN for this interface. If the value of VirtualWwnAuto is 'true', then value of this virtual pWWN is automatically generated by the device. If value of this pWWN is set explicitly, then value of VirtualWwnAuto is implicitly set to 'false'. If length of pWWN is zero, then automatic virtual WWN generation is disabled. This pWWN can not be set to length zero
Auto	Enable automatic generation of Virtual WWNs on this interface. If the value of VirtualWwnPwwn is set explicitly, then the value of Auto will be implicitly set to false. Also, if this Auto is set to 'true', then value of VirtualWwnPwwn is overwritten with auto generated virtual port WWN.

LastChange	The value of sysUpTime at the time of the last change to this Virtual WWN Entry.
------------	----------------------------------------------------------------------------------

FlexAttach Physical to Virtual WWNs

Field	Description
virtual pWWN	This is the virtual port WWN for this device port WWN. In order to minimize WWN collision, no two instances of this Virtual pWWN can have same value. Note : The Virtual pWWN cannot be changed when corresponding device is logged in.
LastChange	The value of sysUpTime at the time of the last change to this Virtual WWN Entry.

FIPS

Field	Description
ModeActivation	To enable/disable FIPS mode on the device. FIPS 140-2 is a set of security requirements for cryptographic modules and it details the U.S. Government requirements for cryptographic modules. A module will comprise both hardware and software, eg a datacenter switching or routing module. The module is said to be in FIPS enabled mode when a request is recieved to enable the FIPS mode and a set of self-tests are successfully run in response to the request. If the self-tests fail, then an appropriate error is returned

FCIP FICON Configuration

Field	Description
Interface	This is a unique value that identifies the interface on this FCIP device to which this link pertains.
VSAN List Admin	This is the list of VSANs (in the range 1 through 2047) for which Ficon Tape Acceleration is configured. Only VSANs with a cficonVsanEntry of CISCO-FICON-MIB present can be configured for Ficon Tape Acceleration.
VSAN List Oper	This is the list of VSANs (in the range 1 through 2047) for which Ficon Tape Acceleration is operationally "ON".

Port Channels AutoCreate

Field	Description
Channel	The channel group mode of this PortChannel.
Persistent	True if the PortChannel is persistent.

SPAN Sessions

Field	Description
Dest Interface	The Span Destination port interface.
Filter VSAN List	The VSANs that are assigned to this session.
Status Admin	Suspend an active session or activate an inactive session.
Status Oper	The current state of the session.
Description	The description of the session status.
VSAN List	The VSANs that are assigned to this session.
Or Interface (Direction)	The destination port ID to be configured for the session.
Inactive Reason	Description of the reason why this session is not active.

Span Global

Field	Description
MaxQueuedSpanPackets	This field specifies the drop threshold packets for all span sessions. The MaxQueuedSpanPackets field is only available when no session is active.

SPAN Source Interfaces

Field	Description
Interface, Direction	The destination port ID configured for the session, and the direction of traffic.

Port Tracking Dependencies

Field	Description
Linked, Destination Interfaces	The interfaces that are doing the tracking.
VSAN Type	Whether a single VSAN or all VSANs are being tracked.
VSAN ID	If a single VSAN is being tracked, the ID of that VSAN.

Port Tracking Force Shut

Field	Description
Interface	The interface of the port to be configured for the forced-shut mode.
Force Shut	If true, the port is brought down administratively, and you must bring the port up manually. If false, the port is brought down operationally only, and is brought up again as soon as any one of the tracked ports comes up.

Port Guard

Field	Description
Interface	Name of the interface
Enable	Specifies whether an interface can be stopped from changing between up and down states or allowed to change states continuously.
Duration (sec)	Specifies the time duration in which a port is allowed to change states.
Number of Flaps	Specifies the number of times the port can flap in the time specified in the Duration.
Oper	Operational state of the interface.

Bandwidth Reservation: 48-Port 96-Gbps Fibre Channel module

RateMode Config Macro	Description
Dedicated 4 Gbps on the first port of each group and the remaining ports 8 Gbps shared	Allocates a rate mode and admin speed of 4 Gbps on the first port of each group and the remaining ports share 8 Gbps depending on the operational speed of the ports
Dedicated 8 Gbps on the first port of each group and the remaining ports 8 Gbps shared	Allocates a rate mode and admin speed of 8 Gbps on the first port of each group and the remaining ports share 8 Gbps depending on the operational speed of the ports
Shared 8 Gbps on all ports (initial & default settings)	Allocates a rate mode and admin speed of 8 Gbps on all the available ports. This is the default setting.

Bandwidth Reservation: 48-Port 48-Gbps Fibre Channel module

RateMode Config Macro	Description
Dedicated 2 Gbps on the first port of each group and the remaining ports 4 Gbps shared	Allocates a rate mode and admin speed of 2Gbps on the first port of each group and the remaining ports share 4 Gbps depending on the operational speed of the ports
Dedicated 8 Gbps on the first port of each group and the remaining ports 4 Gbps shared	Allocates a rate mode and admin speed of 8 Gbps on the first port of each group and the remaining ports share 4 Gbps depending on the operational speed of the ports

RateMode Config Macro	Description
Shared Auto with Maximum of 4 Gbps on all ports (initial & default settings)	Allocates a maximum rate mode and admin speed of 4Gbps on all the available ports. This is the default setting.

Bandwidth Reservation: 24-Port 48-Gbps Fibre Channel module

RateMode Config Macro	Description
Dedicated 8 Gbps on the first port of each group and the remaining ports 8G shared	Allocates a rate mode and admin speed of 8Gbps on the first port of each group and the remaining ports share 8 Gbps depending on the operational speed of the ports
Shared Auto on all ports (initial & default settings)	Allocates a rate mode and admin speed of 8 Gbps on all the available ports. This is the default setting.

Bandwidth Reservation: 48-Port 256-Gbps Fibre Channel module

RateMode Config Macro	Description
Dedicated 8 Gbps on the first 4 ports in each 6-port port group and the remaining ports 8G shared	Allocates a rate mode and admin speed of 8Gbps on the first 4 ports in each 6-port port group and the remaining ports share 8 Gbps depending on the operational speed of the ports.
Dedicated 8 Gbps on the first port of each group and the remaining ports 8G shared	Allocates a rate mode and admin speed of 8 Gbps on the first port of each group and the remaining ports share 8 Gbps depending on the operational speed of the ports.
Shared 8G 0n all ports	Allocates a rate mode and admin speed of 8 Gbps on all the available ports. This is the default setting.
Dedicated 4G 0n all ports	Allocates a rate mode and admin speed of 4Gbps on all the available ports.
Dedicated 10G on following ports: <ul style="list-style-type: none"> • 4,5,6,7,8,10 (ports 1,2,3,9,11,12 disabled) • 16, 17, 18, 19, 20, 22 (ports 13,14,15, 21,23,24 disabled) • 28,29,30,31,32,34 (ports 25,26,27,33,35,36 disabled) • 40,41,42,43,44,46 (ports 37,38, 39 45, 47, 48 disabled) 	Allocates a rate mode and admin speed of 10Gbps on the following ports.

Bandwidth Reservation: 32-Port 256-Gbps Fibre Channel module

RateMode Config Macro	Description
Dedicated 8 Gbps on on all ports	Allocates a rate mode and admin speed of 8 Gbps on all the available ports.

RateMode Config Macro	Description
Shared 8 Gbps on on all ports — initial and default settings.	Allocates a rate mode and admin speed of shared 8 Gbps on all the available ports.
Dedicated 10G on following ports: <ul style="list-style-type: none"> • 2,3,4,5,6,8 (ports 1 and 7 disabled) • 10,11,12,13,14,16 (ports 9 and 15 disabled) • 18,19,20,21,22,24 (ports 17 and 23 disabled) • 26,27,28,29,30,32 (ports 25 and 31 disabled) 	Allocates a rate mode and admin speed of 10Gbps on the specified ports.

DS-X9448-768K9 (Luke) Line Card Bandwidth Reservation

RateMode Config Macro	Description
Dedicated 10G on the following ports: <ul style="list-style-type: none"> • Ports 1-8 • Ports 9-16 • Ports 17-24 • Ports 25-32 • Ports 33-40 • Ports 41-48 	Allocates dedicated rate mode and admin speed of 10 Gbps on the specified ports.
Unconfigure 10G on the following ports: <ul style="list-style-type: none"> • Ports 1-8 • Ports 9-16 • Ports 17-24 • Ports 25-32 • Ports 33-40 • Ports 41-48 	Reverts to default rate mode and admin speed on the specified ports. Transceiver frequency is set to FC. This operation is disruptive.

FC

This section includes the following:

VSAN General

Field	Description
Name	The name of the VSAN. Note that default value will be the string VSANxxxx where xxxx is value of vsanIndex expressed as 4 digits. For example, if vsanIndex is 23, the default value is VSAN0023.
Mtu	The MTU of the VSAN. Normally, this is 2112.

Field	Description
LoadBalancing	The type of load balancing used on this VSAN. <ul style="list-style-type: none"> • srcdst - use source and destination ID for path selection • srcdst 0xld - use source, destination, and exchange IDs
InterOp	The interoperability mode of the local switch on this VSAN. <ul style="list-style-type: none"> • standard • interop-1 • interop-2 • interop-3
AdminState	The state of this VSAN.
OperState	The operational state of the VSAN.
InOrderDelivery	The InorderDelivery guarantee flag of device. If true, then the inorder delivery is guaranteed. If false, it is not guaranteed.
DomainId	Specifies an insistent domain ID.
FICON	True if the VSAN is FICON-enabled.
Network Latency	Network latency of this switch on this VSAN. This is the time interval after which the frames are dropped if they are not delivered in the order they were transmitted.

VSAN Membership

Field	Description
Switch	Name of the switch
Ports	FC Ports in VSAN
Channels	PortChannels in VSAN
FCIP	FCIP Interfaces in VSAN
iSCSI	iSCSI Interfaces in VSAN
FICON	Interfaces in VSAN by FICON
FC Virtual Interface	Virtual FC interfaces in VSAN

VSAN Interop-4 WWN

Field	Description
VSAN ID	The ID of the VSAN containing the McData switch.
WWN	The WWN of the McData switch.

VSAN Timers

Field	Description
VSAN Id	The ID of the VSAN.
R_A_TOV	The Resource_Allocation_Timeout Value used for FxPorts as the timeout value for determining when to reuse an NxPort resource such as a Recovery_Qualifier. It represents E_D_TOV plus twice the maximum time that a frame may be delayed within the Fabric and still be delivered. Note that all switches in a fabric should be configured with the same value of this timeout.
D_S_TOV	The Distributed_Services_Timeout Value which indicates that how long a distributed services requestor will wait for a response.
E_D_TOV	The Error_Detect_Timeout Value used for FxPorts as the timeout value for detecting an error condition. Note that all switches in a fabric should be configured with the same value of this timeout. Note that value must be less than value of D_S_TOV.
NetworkDropLatency	Network latency of this switch on this VSAN.

VSAN Default Zone Policies

Field	Description
Zone Behavior	Represents the initial value for default zone behavior on a VSAN when it is created. If a VSAN were to be deleted and re-created again, the default zone behavior will be set to the value specified for this object.
Propagation Mode	Represents the initial value for zone set propagation mode on a VSAN when it is created. If a VSAN were to be deleted and re-created again, the zone set propagation mode will be set to the value specified for this object.

IVR Local Topology

Field	Description
VSAN List	The list of configured VSANs that are part of IVR topology on this device.

IVR Fabric ID

Field	Description
VSAN List	The list of configured VSANs that are part of IVR topology on this device.

IVR Default Fabric ID

Field	Description
Fabric Id	The configured Default Autonomous Fabric ID of this switch.

IVR Action

Field	Description
Activate Local Topology	Setting this object to activate is a request for the configured IVR topology to be activated on this device. i.e., for the current configuration of IVR topology to be cloned, with the clone becoming the active IVR topology.
IsActive	This object indicates if IVR topology is active or not. If true, the IVR topology is active. If false, the IVR topology is not active.
Activation Time	When the IVR topology was most recently activated. If the IVR topology has not been activated prior to the last re-initialization of the local network management system, then this value will be N/A.
Enable IVR NAT	Enable FCID and VSAN identifier translation across VSAN boundaries. If true, the VSAN identifier as well as the entire FCID of the end devices would be modified as frames cross VSAN boundaries.
Auto Discover Topology	Enable automatic VSAN topology discovery. If true, automatic VSAN topology discovery is turned on. IVR processes would communicate with each other to provide a global view of the physical topology to all the IVR enabled switches. If false, automatic VSAN topology discovery is turned off.

IVR RDI VSANs

Field	Description
Add Virtual Domain to FC Domain List	This object lists VSANs in which the virtual domains in a VSAN are added to the domain list in that VSAN.

IVR Active Topology

Field	Description
VSAN List	The list of VSANs that are part of IVR topology on this device.

IVR Zoneset Status

Field	Description
Status	Status of the active IV Zoneset on this VSAN.
	<ul style="list-style-type: none"> • idle - Idle • active - Active • deactive - Deactive • defaultZoneDeny - Activation failed because of default zone behavior is deny and there is no regular active zoneset. • activationFailed - Activation failed • deactivationFailed - Deactivation failed • activationNotInitiated - Activation not initiated • activationFailedFabricChgFailed - Activation failed because of fabric change failed. • deactivationNotInitiated - Deactivation not initiated. • deactivationFailedFabChgFailed - Deactivation failed because of fabric change failed. • deactivationNotInitiated - Deactivation not initiated. • deactivationFailedFabChgFailed - Deactivation failed because of fabric changing. • activating - Activation in progress. • activatingWaitForLowestSwwn - Activation in progress; waiting for the lowest switch WWN switch to add IV zoneset to the regular active zoneset. • activatingFabricChanging - Activation in progress; fabric is changing. • deactivating - Deactivation in progress. • deactivatingWaitForLowestSwwn - Deactivation in progress; waiting for the lowest switch WWN switch to delete IVR zoneset from the regular active zoneset. • deactivatingFabricChanging - Deactivation in progress; fabric is changing. • defaultZonePermit - Activation failed because of default zone behavior is permit. • defaultZonePermitNoForce - Activation failed because of default zone behavior is permit with no force option. • defaultZonePermitActZsNoForce - Activation failed because of default zone behavior is permit and with regular activate zoneset and no force option. • denyNoActiveZoneset - Activation failed because there is no active zoneset. • activationFailedLowestWwnWait - Activation failed waiting for the switch with lowest wwn to activate this zoneset. • deactivationFailedLowestWwnWait - Deactivation failed waiting for switch with lowest wwn to deactivate this zoneset. • activationFailedZoneNmCtnsIIChar - Activation fails because one of the zone names in zoneset that is being activated contains illegal character.

IVR Discrepancies

Field	Description
Discrepancy	The checksum of the enforced (active) IV zoneset.
RegionID	Identifies the CFS configuration supported region.

IVR Domains

Field	Description
Domain Id	The FC domain ID that will be used to represent the VSAN.

IVR FCID

Field	Description
FCID	The FCID to be used by IVR to represent the device.

IVR Zoneset Active Zones

Field	Description
VSAN Id	IVR VSAN ID.
Zone	Active IVR zone name.
Fabric Id	Autonomous fabric ID.
Switch Interface	Switch interface to which the zone member is connected to.
Name	Zone member name.
WWN	Zone member WWN.
FcId	Zone member FC ID.
LUNs	Zone member LUN.
Status	<ul style="list-style-type: none"> • Not in Fabric: If zone member is not in the fabric. • Not in VSAN: If zone member is not present in the VSAN. • n/a: Cannot determine status. <p>Empty: Member is present in fabric and correct VSAN and can communicate with other members of the zone.</p>

IVR Zoneset Active Zones Attributes

Field	Description
Zone	Active IVR zone name.
QoS	True if QoS enabled, otherwise false.
QoS Priority	QoS priority value (Low, Medium, or High).
Broadcast	Specifies if broadcast zoning is enabled on this default zone on this VSAN. If true, then it is enabled. If false, then it is disabled.

IVR Zoneset Name

Field	Description
VSAN Id	IVR VSAN ID.
Zone	Active IVR zone name.
Fabric Id	Autonomous fabric ID.
Switch Interface	Switch interface to which the zone member is connected to.
Name	Zone member name.
WWN	Zone member WWN.
FcId	Zone member FC ID.
Luns	Zone member LUN.
Status	<ul style="list-style-type: none"> • Not in Fabric: If zone member is not in the fabric. • Not in VSAN: If zone member is not present in the VSAN. • n/a: Cannot determine status. <p>Empty: Member is present in fabric and correct VSAN and can communicate with other members of the zone.</p>

DPVM Actions

Field	Description
Action	Helps in activating the set of bindings.
Result	Indicates the outcome of the activation.
Status	Indicates the state of activation. If true, then activation has been attempted as the most recent operation. If false, then an activation has not been attempted as the most recent operation.
CopyActive to Config	When set to copy(1), results in the active (enforced) binding database to be copied on to the configuration binding database. The learned entries are also copied.

Field	Description
Auto Learn Enable	Helps to learn the configuration of devices logged into the local device on all its ports and the VSANs to which they are associated.
Auto Learn Clear	Assists in clearing the auto-learned entries.
Clear WWN	Represents the Port WWN (pWWN) to be used for clearing its corresponding auto-learned entry.

DPVM Config Database

Field	Description
Switch	Name of the switch.
Type	Specifies the type of the corresponding instance of device.
WWN or Name or MAC	Represents the logging in device. The value depends on the corresponding device type (PWWN, NWWN or MAC).
VSAN Id	Represents the VSAN to be associated to the port on the local device on which the device represented by cdpvmLoginDev logs in.
Switch Interface	Represents the device alias.

DPVM Active Database

Field	Description
Type	Specifies the type of the corresponding instance of cdpvmEnfLoginDev.
WWN or Name or MAC	Represents the logging in device. The value depends on the corresponding device type (PWWN, NWWN or MAC).
VSAN Id	Represents the VSAN of the port on the local device through which the device represented by cdpvmEnfLoginDev logs in.
Interface	Represents the device alias.
IsLearnt	Indicates whether this is a learnt entry or not. If true, then it is a learnt entry. If false, then it is not.

Domain Manager Running

Field	Description
State	The state of the Domain Manager on the local switch on this VSAN.
DomainId	The Domain ID of the local switch on this VSAN or 0 if no Domain ID has been assigned.
Local Switch WWN	The WWN of the local switch on this VSAN.
Local Priority	The running priority of the local switch on this VSAN.

Field	Description
Principal Switch WWN	The WWN of the principal switch on this VSAN, or empty string if the identity of the principal switch is unknown.
Principal Priority	The running priority of the principal switch on this VSAN.

Domain Manager Configuration

Field	Description
Enable	Enables the Domain Manager on this VSAN. If enabled on an active VSAN, the switch will participate in principal switch selection. If disabled, the switch will participate in neither the principal switch selection nor domain allocation. Thus, Domain ID needs to be configured statically.
Running DomainId	<p>The configured Domain ID of the local switch on this VSAN or 0 if no Domain ID has been configured. The meaning depends on DomainIdType.</p> <p>If Type is 'preferred', then domain ID configured is called 'preferred Domain ID'. The valid values are between 0 and 239. In a situation where this domain could not be assigned, any domain ID would be acceptable. The value '0' means any domain ID.</p> <p>If Type is 'static' (insistent), then domain ID is called 'static Domain ID' and valid values are between 1 and 239. In a situation where this domain was non-zero but could not be assigned, no other domain ID would be acceptable.</p> <p>If the Domain Manager is enabled on the VSAN, then a RDI (Request Domain ID) will be sent requesting this Domain ID. If no Domain ID can be granted in the case of 'preferred' or if the configured 'static' (insistent) domain ID cannot be not granted then, it is an error condition. When this error occurs, the E_ports on that VSAN will be isolated.</p> <p>If the domain manager is not enabled, then the static (insistent) Domain ID is assumed to be granted, if it has been configured (to a valid number). If either of the domain IDs are not configured with a non-zero value on this VSAN and if the domain manager is not enabled, then - switch will isolate all of its E_ports on this VSAN.</p>
DomainId Type	Type of configured Domain ID.
FabricName	The WWN that is used for fabric logins on this VSAN. This is used only if Enable is false. If Enable is true, then principal switch WWN is used. It is automatically set to the default value when set to zero-length value.
Priority	Priority of the switch to be used in principal switch selection process.
Contiguous Allocation	Determines how the switch behaves when elected as the principal switch. If true, switch won't accept non-contiguous domain IDs in RDIs and will try to replace all the Domain IDs in the list with contiguous domain IDs if a RDI for a contiguous Domain ID can not be fulfilled. If false, then the switch acts normally in granting the Domain IDs even if they are not contiguous.
Auto Reconfigure	Determines how the switch responds to certain error conditions. The condition that can cause these errors is merging of two disjoint fabrics that have overlapping Domain ID list. If true, the switch will send a RCF (ReConfigureFabric) to rebuild the Fabric. If false, the switch will isolate the E_ports on which the errors happened.

Field	Description
Persistent FcId	If true, then all the FC ID assigned on this VSAN are made persistent on this VSAN. If false, then all the entries on VSAN in PersistencyTable are deleted.
Purge FcIds?	Tells the Domain Manager to purge the FC IDs on this VSAN in the FC ID persistency database.
Restart?	Tells the Domain Manager to rebuild the Domain ID tree all over again. If 'disruptive', then a RCF (ReConfigure Fabric) is generated in the VSAN in order for the fabric to recover from the errors. If nonDisruptive, then a BF (Build Fabric) is generated in the VSAN.
Optimization	You need to click the field to select one of the following. To disable turbo mode, do not select anything. <ul style="list-style-type: none"> • Fast-Restart- Set the optimization type to fast restart. • Selective-Restart- Set the optimization type to selective restart.

Domain Manager Domains

Field	Description
SwitchWWN	The WWN of the switch to which the corresponding value of DomainId is currently assigned for the particular VSAN.

Domain Manager Statistics

Field	Description
Prin. Sel Total	The number of principal switch selections on this VSAN.
Prin. Sel Local	The number of times the local switch became the principal switch on this VSAN.
Fabric Builds (BF)	The number of BuildFabrics (BFs) that have occurred on this VSAN.
Fabric Reconfigures (Rcf)	The number of ReconfigureFabrics (RCFs) that have occurred on this VSAN.
FcIds Free	The number of FC IDs that are unassigned on this VSAN.
FcIds Assigned	The number of FC IDs that are assigned on this VSAN.
FcIds Reserved	The number of FC IDs that are reserved on this VSAN.

Domain Manager Interfaces

Field	Description
Role	One of the following: <ul style="list-style-type: none"> • nonPrincipal - non-principal interface • principalUpstream - upstream principal interface • principalDownstream - downstream principal interface • isolated - isolated interface • down - down interface unknown • unknown - unknown interface
RcfReject	Determines if the incoming ReConfigure Fabric (RCF) messages on this interface on this VSAN is accepted or not. If true, then the incoming RCF is rejected. If false, incoming RCF is accepted. Note that this does not apply to the outgoing RCFs generated by this interface.

Domain Manager Persistent Fclds

Field	Description
FcId	The FC ID assigned for this WWN on this VSAN. The third octet must be 0x00 if value of PersistencyNum is area.
Mask	The number of FC IDs starting from PersistencyFcId which are assigned either statically or dynamically for this WWN on this VSAN. The value one means just one FC ID is assigned. The value area means all the FC IDs in the area that is specified in the second octet of FcId are assigned. Typically, 256 FC IDs are assigned for an area. This value cannot be changed if the value of Used is true.
Used	Indicates if this FC ID is used or not.
Assignment	The type of persistency of this FC ID.

Domain Manager Allowed DomainIds

Field	Description
List	Provides the lists of domains that are allowed. A domain is allowed in this VSAN if the corresponding bit has a value of 1. If it has a value which is less than 32 bytes long, then the domains which are not represented are not considered to be in the list. If this object is a zero-length string, then no domains are allowed in this VSAN.

Zoneset Active Zones

Field	Description
Zone	Zone name.

Field	Description
Type	Zone member type.
Switch Interface	Switch interface to which the zone member is connected to.
Name	Zone member name.
WWN	Zone member WWN.
FcId	Zone member FC ID.
LUNs	Zone member LUN.
Status	<ul style="list-style-type: none"> • Not in Fabric: If zone member is not in the fabric. • Not in VSAN: If zone member is not present in the VSAN. • n/a: Cannot determine status. • Empty: Member is present in fabric and correct VSAN and can communicate with other members of the zone.

Zoneset Unzoned

Field	Description
Name	Zone member name.
WWN	Zone member WWN.
FcId	Zone member FC ID.

Zoneset Status

Field	Description
Status	Indicates the outcome of the most recent activation/deactivation.
Activation Time	When this entry was most recently activated. If this entry has been activated prior to the last re-initialization of the local network management system, then this value will be N/A.
FailureCause	The reason for the failure of the zoneset activation/deactivation.
FailedSwitch	The domain ID of the device in the fabric that has caused the Change Protocol to fail.
Active == Local?	Indicates whether the enforced database is the same as the local database on this VSAN. If true, then they are the same. If false, then they are not the same.
Active Zoneset	The name of the enforced IV zoneset.
Hard Zoning	Indicates whether the hard zoning is enabled on this VSAN. Hard zoning is a mechanism by which zoning is enforced in hardware. If true, then hard zoning is enabled on this VSAN. If false, then hard zoning is not enabled on this VSAN.

Zoneset Policies

Field	Description
Default Zone Behavior	Controls the behavior of the default zone on this VSAN. If it is set to permit, then the members of the default zone on this VSAN can communicate with each other. If it is set to deny, then the members of the default zone on this VSAN cannot communicate with each other.
Default Zone ReadOnly	Indicates whether SCSI read operations are allowed on members of the default zone which are SCSI targets, on this VSAN. If true, then only SCSI read operations are permitted. So, this default zone becomes a read-only default zone on this VSAN. If false, then both SCSI read and write operations are permitted.
Default Zone QoS	Specifies whether the QoS attribute for the default zone on this VSAN is enabled. If true, then QoS attribute for the default zone on this VSAN is enabled. If false, then the QoS attribute for the default zone on this VSAN is disabled.
Default Zone QoS Priority	Specifies the QoS priority value.
Default Zone Broadcast	Specifies if broadcast zoning is enabled on this default zone on this VSAN. If true, then it is enabled. If false, then it is disabled.
Propagation	Controls the way zoneset information is propagated during Merge/Change protocols on this VSAN
Read From	Specifies whether the management station wishes to read from the effective database or from the copy database.

Zoneset Active Zones Attributes

Field	Description
Name	Zone name.
Read Only	Indicates if only SCSI read operations are allowed on members of the default zone which are SCSI targets on this VSAN. If true, then only SCSI read operations are permitted. So, this default zone becomes a read-only default zone on this VSAN. If false, then both SCSI read and write operations are permitted.
QoS	Specifies whether the QoS attribute for the default zone on this VSAN is enabled. If true, then QoS attribute for the default zone on this VSAN is enabled. If false, then the QoS attribute for the default zone on this VSAN is disabled.
QoS Priority	Specifies QoS priority value (Low, Medium, or High).
Broadcast	Specifies if broadcast zoning is enabled on this default zone on this VSAN. If true, then it is enabled. If false, then it is disabled.

Zoneset Enhanced

Field	Description
Action	When set to basic(1), results in the zone server operating in the basic mode as defined by FC-GS4 standards. When set to enhanced(2), results in the zone server operating in the enhanced mode as defined by FC-GS4 standards.
Result	The outcome of setting the mode of operation of the local Zone Server on this VSAN.
Config DB Locked By	Specifies the owner for this session.
Config DB Discard Changes	Assists in committing or clearing the contents of the copy database on this session.
Config DB Result	Indicates the outcome of setting the corresponding instance of czseSessionCntl to commitChanges(1).
Enforce Full DB Merge	Controls the zone merge behavior. If this object is set to allow, then the merge takes place according to the merge rules. If set to restrict, then if the merging databases are not exactly identical, the Inter-Switch Link (ISL) between the devices is isolated.
Read From	Specifies whether the management station wishes to read from the effective database or from the copy database.

Zoneset Read Only Violations

Field	Description
Violations	The number of Data protected Check Condition error responses sent by the local Zone Server.

Zoneset Statistics

Field	Description
Merge Req Tx	The number of Merge Request Frames sent by this Zone Server to other Zone Servers in the fabric on this VSAN.
Merge Req Rx	The number of Merge Request Frames received by this Zone Server from other Zone Servers in the fabric on this VSAN.
Merge Acc Tx	The number of Merge Accept Frames sent by this Zone Server to other Zone Servers in the fabric on this VSAN.
Merge Acc Rx	The number of Merge Accept Frames received by this Zone Server from other Zone Servers in the fabric on this VSAN.
Change Req Tx	The number of Change Requests sent by this Zone Server to other Zone Servers in the fabric on this VSAN.
Change Req Rx	The number of Change Requests received by this Zone Server from other Zone Servers in the fabric on this VSAN.

Field	Description
Change Acc Tx	The number of Change Responses sent by this Zone Server to other Zone Servers in the fabric on this VSAN.
Change Acc Rx	The number of Change Responses received by this Zone Server from other Zone Servers in the fabric on this VSAN.
GS3 Rej Tx	The number of GS3 requests rejected by this Zone Server on this VSAN.
GS3 Req Rx	The number of GS3 requests received by this Zone Server on this VSAN.

Zoneset LUN Zoning Statistics

Field	Description
INQUIRY	The number of SCSI INQUIRY commands that have been received by the local zone server.
REPORT LUN	The number of SCSI Report LUNs commands that have been received by the local zone server. Typically the Report LUNs command is sent only for LUN 0.
SENSE	The number of SCSI SENSE commands that have been received by the local zone server.
Other Cmds	The number of SCSI Read, Write, Seek, etc., commands received by the local zone server.
BadInquiry Errors	The number of No LU error responses sent by the local zone server.
Illegal Errors	The number of Illegal Request Check Condition responses sent by the local zone server.

Zoneset Members

Field	Description
Zone	Default zone.
Type	FCID.
Switch Interface	Switch interface to which the zone member is connected to.
Name	Zone member name.
WWN	Zone member WWN.
FcId	Zone member FC ID.
Luns	Zone member LUN.
Status	<ul style="list-style-type: none"> • Not in Fabric: If zone member is not in the fabric. • Not in VSAN: If zone member is not present in the VSAN. • n/a: Cannot determine status. <p>Empty: Member is present in fabric and correct VSAN and can communicate with other members of the zone.</p>

Fabric Config Server Discovery

Field	Description
Status	<p>The status of the discovery on the local switch. Initially when the switch comes up, this will be set to databaseInvalid state on all VSANs. This indicates that a discovery needs to be done. The state will be set to inProgress for this VSAN during the discovery. Once the discovery is completed on this VSAN, this will be set to completed. After the discovery is completed for the specified list of VSANs, the data is cached for an interval of time.</p> <p>Once this interval of time expires, the data is lost and this will be set to databaseInvalid state for the specified list of VSANs.</p>
CompleteTime	When the last discovery was completed on this VSAN. This value is N/A before the first discovery on this VSAN.

Fabric Config Server Interconnect Elements

Field	Description
Type	The type of this Interconnect Element.
DomainId	The Domain Id of this Interconnect Element. If the Domain Id has not been configured, then this value is 0.
MgmtId	The management identifier of this Interconnect Element. If the Interconnect Element is a switch, then this will be the Domain Controller identifier of the switch.
FabricName	The fabric name of this Interconnect Element.
LogicalName	The logical name of this Interconnect Element.
Vendor, Model, Release, WWN	The information list corresponding to this Interconnect Element.
MgmtAddrList	The management address list corresponding to this Interconnect Element.

Fabric Config Server Platforms (Enclosures)

Field	Description
Name	The name of this platform.
Type	The type of this platform.
ConfigSource	The source of configuration of this entry. Note that an entry which is configured via GS3 cannot be deleted through SNMP.
NodeList	The node name list corresponding to this platform.

Field	Description
MgmtAddrList	The management address list corresponding to this Platform.

Fabric Config Server Fabric Ports

Field	Description
Type	The type of this port.
TXType	The TX type of this port.
ModuleType	The module type of this port.
Interface	The physical number corresponding to this port entry.
State	The state of this port.
AttachedPortList	The attached port name list corresponding to this port.

FC Routes

Field	Description
Preference	The value used to select one route over another when more than one route to the same destination is learned from different protocols, peers, or static routes. The preference value is an arbitrarily assigned value used to determine the order of routes to the same destination in a single routing database (RIB). The active route is chosen by the lowest preference value.
LastChangeTime	The last time a row was created, modified, or deleted in the FC route table.
DomainId	The domain ID of next hop switch. However, when read, this value could be N/A if the value of fcRouteProto is local.
Metric	The routing metric for this route. The use is dependent on fcRouteProto used.
Type	The type of route. <ul style="list-style-type: none"> • local(1): refers to a route for which the next hop is the final destination. • remote(2): refers to a route for which the next hop is not the final destination. This is not relevant for multicast and broadcast route entries.

FDMI HBAs

Field	Description
Sn	The serial number of this HBA.
Model	The model of this HBA.
ModelDescr	The model description.
OSInfo	The type and version of the operating system controlling this HBA.
MaxCTPayload	The maximum size of the Common Transport (CT) payload including all CT headers but no FC frame header(s), that may be send or received by application software resident in the host containing this HBA.

FDMI Ports

Field	Description
SupportedFC4Type	The supported FC-4 types attribute registered for this port on this VSAN.
SupportedSpeed	The supported speed registered for this port on this VSAN.
CurrentSpeed	The current speed registered for this port on this VSAN.
MaxFrameSize	The maximum frame size attribute registered for this port on this VSAN.
OsDevName	The OS Device Name attribute registered for this port on this VSAN.
HostName	The name of the host associated with this port.

FDMI Versions

Field	Description
Hardware	The hardware version of this HBA.
DriverVer	The version level of the driver software controlling this HBA.
OptROMVer	The version of the Option ROM or the BIOS of this HBA.
Firmware	The version of the firmware executed by this HBA.

Flow Statistics

Field	Description
Type	The matching criteria by which flows are selected to be included in the traffic which is instrumented by the ingress traffic counters.
VsanId	The id of VSAN.
DestId	The destination fibre channel address ID.
SrcId	The source fibre channel address ID.
Mask	The mask for source and destination fibre channel address ID.
Frames	The number of received frames for the flow created by the network manager.
Bytes	The number of received frame bytes for the flow created by the network manager.
CreationTime	The timestamp indicating the time the row was created or modified.

FCC

Field	Description
Enable	Enable Fabric Congestion Control
Priority	Specifies the priority level for the frames.
EdgeQuenchPktsRecd	The number of Edge Quench packets received and processed on this port.
EdgeQuenchPktsSent	The number of Edge Quench packets generated on this Port as result of congestion.
PathQuenchPktsRecd	The number of Path Quench packets received and processed on this port.
PathQuenchPktsSent	The number of Path Quench packets generated on this Port as result of congestion.
CurrentCongestionState	The current FCC congestion state of this Port indicating the severity of the congestion.
LastCongestedTime	When the congestion state of the Port changed to noCongestion from some other value. N/A if the congestion state of the Port has never transitioned to noCongestion since the last restart of the device.
LastCongestionStartTime	When the congestion state of the port changed from noCongestion to some other value.

Field	Description
IsRateLimitingApplied	If true, rate limiting is currently being applied on this port.

Diagnostics

Field	Description
Value	Displays the most recent measurement seen by the sensor.
Alarms High and Low	Represents the severity level of the SFP diagnostic information of an interface for temperature, voltage, current, optical transmit and receive power. It ranges from 1 to 6, with 6 being highest severity.
Warnings High and Low	

FSPF General

Field	Description
AdminStatus	The desired state of FSPF on this VSAN.
OperStatus	State of FSPF on this VSAN.
SetToDefault	Enabling this changes each value in this row to its default value. If all the configuration parameters have their default values and if the VSAN is suspended, then the row is deleted automatically.
RegionId	The autonomous region of the local switch on this VSAN.
DomainId	The Domain Id of the local switch on this VSAN.
SpfHoldTime	The minimum time between two consecutive SPF computations on this VSAN. The smaller value means that routing will react to the changes faster but the CPU usage is greater.
SpfDelay	The time between when FSPF receives topology updates and when it starts the Shortest Path First (SPF) computation on this VSAN. The smaller value means that routing will react to the changes faster but the CPU usage is greater.
MinLsArrival	The minimum time after accepting a Link State Record (LSR) on this VSAN before accepting another update of the same LSR on the same VSAN. An LSR update that is not accepted because of this time interval is discarded.

Field	Description
MinLsInterval	The minimum time after this switch sends an LSR on this VSAN before it will send another update of the same LSR on the same VSAN.
LsRefreshTime	The interval between transmission of refresh LSRs on this VSAN.
LSRMaxAge	The maximum age an LSR will be retained in the FSPF database on this VSAN. It is removed from the database after MaxAge is reached.
CreateTime	When this entry was last created.
Checksum	The total checksum of all the LSRs on this VSAN.

FSPF Interfaces

Field	Description
SetToDefault	Enabling this changes each value in this row to its default value. If all the configuration parameters have their default values and if the interface is down, then the row is deleted automatically.
Cost	<p>The administrative cost of sending a frame on this interface on this VSAN. The value 0 means that the cost has not been configured. Once the value has been configured, the value can not again be 0; so, obviously the value can not be set to 0. If the value is 0 and the corresponding interface is up, the agent sets a value calculated using the ifSpeed of the interface. Otherwise, the value is used as the cost.</p> <p>Note that following formula is used to calculate the link cost.</p> <p>Link Cost = { fspfIfCost if fspfIfCost > 0 {(1.0625e12 / Baud Rate) if fspfIfCost == 0 where Baud Rate is the ifSpeed of the interface.</p>
AdminStatus	The desired state of FSPF on this interface on this VSAN.
HelloInterval	Interval between the periodic HELLO messages sent on this interface on this VSAN to verify the link health. Note that this value must be same on both the interfaces on each end of the link on this VSAN.

Field	Description
DeadInterval	Maximum time for which no HELLO messages can be received on this interface on this VSAN. After this time, the interface is assumed to be broken and removed from the database. Note that this value must be greater than the HELLO interval specified on this interface on this VSAN.
RetransmitInterval	Time after which an unacknowledged link update is retransmitted on this interface on this VSAN.
Neighbour State	The state of FSPF's neighbor state machine, which is the operational state of the interaction with the neighbor's interface which is connected to this interface.
Neighbour DomainId	The Domain ID of the neighbor on this VSAN.
Neighbour PortIndex	The index, as known by the neighbor, of the neighbor's interface which is connected to this interface on this VSAN.
CreateTime	When this entry was last created.

FSPF Interface Stats

Field	Description
CreateTime	When this entry was last created.
ErrorRxPkts	Number of invalid FSPF control frames received on this interface on this VSAN since the creation of the entry.
InactivityExpirations	Number of times the inactivity timer has expired on this interface on this VSAN since the creation of the entry.
LsuRxPkts	Number of Link State Update (LSU) frames received on this interface on this VSAN since the creation of the entry.
LsuTxPkts	Number of Link State Update (LSU) frames transmitted on this interface on this VSAN since the creation of the entry.
RetransmittedLsuTxPkts	Number of LSU frames retransmitted on this interface on this VSAN since the creation of the entry.
LsaRxPkts	Number of Link State Acknowledgement (LSA) frames received on this interface on this VSAN since the creation of the entry.

Field	Description
LsaTxPkts	Number of Link State Acknowledgement (LSA) frames transmitted on this interface on this VSAN since the creation of the entry.
HelloTxPkts	Number of HELLO frames transmitted on this interface on this VSAN since the creation of the entry.
HelloRxPkts	Number of HELLO frames received on this interface on this VSAN since the creation of the entry.

SDV Virtual Devices

Field	Description
Name	Represents the name of this virtual device.
Virtual Domain	The user preference for a persistent Domain ID for this virtual device to indicate a specific partition (domain) of the fabric that this virtual device should belong to.
Virtual FCID	The user preference for a persistent FCID for this virtual device.
Port WWN	The assigned PWWN for this virtual device. The agent assigns this value when the configuration is committed.
Node WWN	The assigned NWWN for this virtual device. The agent assigns this value when the configuration is committed.
Assigned FCID	The assigned FCID of this virtual device. The agent assigns this value when the configuration is committed and the real device that this virtual device virtualizes is on-line.
Real Device Map List	The set of real device(s) that this virtual device virtualizes in this VSAN.

SDV Real Devices

Field	Description
Type	The type of real device identifier represented by the value of the corresponding instance of cFcSdvVirtRealDeviceId that this virtual device virtualizes to.
Name	Represents a real device(s) identifier that this virtual device virtualizes.

Field	Description
Map Type	The mapping association type of the real device(s) (initiator/target).

LUN Discover

Field	Description
StartDiscovery	If Local, then only the directly attached SCSI target devices/ports and LUNs associated with them on all VSANs will be discovered. If Remote, then all SCSI target devices/ports and LUNs associated with them on all VSANs in the whole fabric, except the directly attached ones, will be discovered.
Type	Selecting targets results in only targets being discovered, without the NS results in both targets and LUNs being discovered.
OS	Specifies the operating system on which the LUNs need to be discovered.
Status	Indicates the outcome of the LUN discovery on the local switch. Contains the status of the most recent discovery. <ul style="list-style-type: none"> • inProgress(1) - indicates that the discovery is still in progress. • completed(2) - indicates that the discovery is complete. • failure(3) - indicates that the discovery encountered a failure.
CompleteTime	When the last discovery was completed. The value will be zero or N/A, if discovery has not been performed since the last system restart.

LUN Targets

Field	Description
VsanId	The VSAN to which this target belongs to.
Port WWN	The name of this authorized/discovered target device or port.
DevType	The device type of the SCSI target.
VendorId	The vendor Id of the SCSI target.
ProductId	The product Id of the SCSI target.
RevLevel	The product revision level of the SCSI target.

Field	Description
OtherInfo	The bytes from 0 to 7 in the INQUIRY command response data.

LUNs

Field	Description
Id	The number of this LUN.
Capacity (MB)	The capacity of this LUN.
SerialNum	The serial number of this LUN.
OS	The operating system for which this LUN was discovered.
FC ID	The Fibre Channel ID for this LUN.

Device Alias

Field	Description
Alias	The device alias of this entry. A device can have only one alias configured.
WWN	The Fibre Channel device which is given a device alias.

Device Alias Configuration

Field	Description
Device Alias	The device alias of this entry. A device can have only one alias configured.
WWN	The Fibre Channel device which is given a device alias.

Device Alias Mode

Field	Description
ConfigMode	Specifies the mode in which the device aliases can be configured. When it is set to basic, the device aliases operate in basic mode of operation. When basic mode is turned on, all MIBs which are using device aliases should internally convert them to their equivalent pWWNs and use the pWWNs. The mechanism to be followed for this conversion is implementation specific. When it is set to enhanced, the Device aliases operate in enhanced mode of operation. When enhanced mode is turned on, all MIBs which are using device aliases should use them as is without any conversion. Since the device aliases are used directly without any conversion, this is the native mode of operation of device aliases.

Device Alias Discrepancies

Field	Description
Discrepancy	Represents the checksum computed over the database represented by cfdaConfigTable and the cfdaConfigMode object. This object is used by a network manager to check if the above mentioned objects have changed on the local device. The method used to compute the checksum is implementation specific.

Name Server General

Field	Description
VSAN Id / FcId	The ID of the VSAN or FC.
Type	The port type of this port.
PortName	The fibre channel Port_Name (WWN) of this Nx_port.
NodeName	The fibre channel Node_Name (WWN) of this Nx_port.
FC4Type/Features	The FC-4 Features associated with this port and the FC-4 Type. Refer to FC-GS3 specification for the format.
FC4 Features	The FC-4 Features associated with this port.
ProcAssoc	The Fibre Channel initial process associator.

Field	Description
FabricPortName	The Fabric Port Name (WWN) of the Fx_port to which this Nx_port is attached.

Name Server Advanced

Field	Description
ClassOfSvc	The class of service indicator.
PortIpAddress	Contains the IP address of the associated port.
NodeIpAddress	The IP address of the node of this Nx_port, as indicated by the Nx_Port in a GS3 message that it transmitted.
SymbolicPortName	The user-defined name of this port.
SymbolicNodeName	The user-defined name of the node of this port.
HardAddress	Extended Link Service (FC-PH-2). Hard Address is the 24-bit NL_Port identifier which consists of - the 8-bit Domain Id in the most significant byte - the 8-bit Area Id in the next most significant byte - the 8-bit AL-PA(Arbitrated Loop Physical Address) which an NL_port attempts acquire during FC-AL initialization in the least significant byte. If the port is not an NL_Port, or if it is an NL_Port but does not have a hard address, then all bits are reported as 0s.
ProcAssoc	The Fibre Channel initial process associator (IPA).
PermanentPortName	The Permanent Port Name of this Nx port. If multiple port names are associated with this Nx port via FDISC (Discover F Port Service Parameters), the Permanent Port Name is the original port name associated with this Nx port at login.

Name Server Proxy

Field	Description
PortName	Name of the proxy port which can register/de-register for other ports on this VSAN. Users can enable third party registrations by setting this value.

Name Server Statistics

Field	Description
Queries Rx	The total number of Get Requests received by the local switch on this VSAN.
Queries Tx	The total number of Get Requests sent by the local switch on this VSAN.
Requests Rx Reg	The total number of Registration Requests received by the local switch on this VSAN.
Requests Rx DeReg	The total number of De-registration Requests received by the local switch on this VSAN.
RSCN Rx	The total number of RSCN commands received by the local switch on this VSAN.
RSCN Tx	The total number of RSCN commands sent by the local switch on this VSAN.
Rejects Tx	The total number of requests rejected by the local switch on this VSAN.

Preferred Path Maps and Routes

Field	Description
VSAN Id, Route Id	The VSAN ID of this FC route map. An arbitrary integer value that identifies a route in this FC route map.
Map Active	Allows the activation/de-activation of all the routes within an FC route map. If true, then all the routes within this FC route map will be activated. If false, then all routes within this FC route map will be de-activated.
Route Strict Preference	Allows changes to the way the preferred path selection logic will select the preferred path. Setting it to true makes the preferred path to select the outgoing interface strictly based on the preference set using the cPrefPathRMapSetIntfPref. When it is set to false, then the preferred path selection logic only performs selection only when the current outgoing interface goes down.
Route Active	Allows the activation/de-activation of the route within an FC route map. If true, then the route will be activated. If false, then the route will be de-activated.
RouteActive	Allows the activation/de-activation of the route within an FC route map. If true, then the route will be activated. If false, then the route will be de-activated.

Preferred Path Maps Active

Field	Description
VSAN Id	The VSAN ID of this FC route map.
GlobalActive	Allows the activation/de-activation of all the routes within an FC route map.

Preferred Path All Match Criteria

Field	Description
VSAN Id, Route Id	The VSAN ID of this FC route map. An arbitrary integer value that identifies a route in this FC route map.
Source FcId	The FC ID that needs to be matched with a source address in a frame for flow classification.
Source Information	Represents the mask associated with the source address.
Source Serial Number	Represents the source serial number.
Source Unit Type	The unit type of the source.
Source Tag	Unique identifier for the source address.
Dest FcId	The FC ID that needs to be matched with a destination address in a frame for flow classification.
Dest Information	Represents the mask associated with the destination address.
Dest Serial Number	Represents the destination serial number.
Dest Unit Type	The unit type of the destination.
Dest Tag	Unique identifier for the destination address.

Preferred Path Active Match Criteria

Field	Description
VSAN Id, Route Id	The VSAN ID of this FC route map. An arbitrary integer value that identifies a route in this FC route map.
Source FcId	The FC ID that needs to be matched with a source address in a frame for flow classification.
Source Information	Represents the mask associated with the source address.
Source Serial Number	Represents the source serial number

Field	Description
Source Unit Type	The unit type of the source.
Source Tag	Unique identifier for the source address.
Dest FcId	The FC ID that needs to be matched with a destination address in a frame for flow classification.
Dest Information	Represents the mask associated with the destination address.
Dest Serial Number	Represents the destination serial number.
Dest Unit Type	The unit type of the source.
Dest Tag	Unique identifier for the destination address.

Preferred Path All Sets

Field	Description
VSAN Id, Route Id, Preference	The VSAN ID of this FC route map. An arbitrary integer value that identifies a route in this FC route map. Preference level, which indicates the metric or cost of the preferred path. The lower the number the higher the preference.
Interface	Represents an interface on the local device on which the matched or classified frame will be forwarded.
IVR Nexthop VSAN	Represents the IVR next hop VSAN ID.

RSCN Nx Registrations

Field	Description
RegType	Indicates the type of registration desired by the subscriber. <ul style="list-style-type: none"> 'fromFabricCtrlr' indicates RSCNs generated by the Fabric Controller. 'fromNxPort' indicates RSCNs generated by Nx_Ports. 'fromBoth' indicates RSCNs generated by Fabric Controller and Nx_Ports.

RSCN Multi-PID Support

Field	Description
Enable	Specifies whether the multi-pid option is enabled on this VSAN.

RSCN Event

Field	Description
TimeOut (msec)	The time (in seconds) before the RSCN event times out.

RSCN Statistics

Field	Description
SCR Rx	The number of SCRs received from Nx_Ports on this VSAN.
SCR RJT	The number of SCR rejected on this VSAN.
RSCN Rx	The number of RSCNs from Nx_Ports received on this VSAN.
RSCN Tx	The total number of RSCNs transmitted on this VSAN.
RSCN RJT	The number of RSCN requests rejected on this VSAN.
SW-RSCN Rx	The number of Inter-Switch Registered State Change Notifications (SW_RSCN) received on this VSAN from other switches.
SW-RSCN Tx	The number of Inter-Switch Registered State Change Notifications (SW_RSCN) transmitted on this VSAN to other switches.
SW-RSCN RJT	The number of SW_RSCN requests rejected on this VSAN.

Multicast Root

Field	Description
DomainId	The domain ID of the multicast root on this VSAN.
ConfigMode	The configured multicast root mode on this VSAN.
OperMode	The operational multicast root mode on this VSAN.

QoS Policy Maps

Field	Description
Name	The name of this classifier entry. The name should be unique.

QoS Class Maps

Field	Description
Name	The name of this filter entry. The name should be unique.
Match	Specifies how the filter should be applied. If true, then all the match statements associated with this filter must be satisfied in order for this filter match to be considered successful. If false, then even if any one of the criteria associated with this filter is satisfied, then the filter match is considered successful.

QoS Match Statements

Field	Description
SrcAddr	An FC address that needs to be matched with the source address in a FC frame.
DstAddr	An FC address that needs to be matched with the destination address in a FC frame.
Interface	An FC interface on the local device on which a frame should arrive in order to be classified by this filter. A value of zero indicates that no interface is configured.
Wildcard	Specifies whether the wild-card option has been set. If true, then the wild-card option is set and all the FC traffic will be considered to match the corresponding multi-field classifier. If false, then the wild-card option is not set.

QoS Class Maps by Policy Maps

Field	Description
Class Map ID	Identifies a Fibre Channel filter.
Priority	Specifies priority value.

QoS Policy Maps by VSAN

Field	Description
VSAN Id, Direction	Specifies the direction of traffic flow on this VSAN.
Policy Map Id	Selects the first Differentiated Services Classifier Element to handle traffic on this VSAN.

QoS DWRR

Field	Description
Weight	The weight associated with this queue.

QoS Rate Limit

Field	Description
Percent	Specifies the rate-limit factor on this interface.

Timers and Policies

Field	Description
R_A_TOV	The Resource_Allocation_Timeout Value used for FxPorts as the timeout value for determining when to reuse an NxPort resource such as a Recovery_Qualifier.
D_S_TOV	The Distributed_Services_Timeout Value which indicates how long a distributed services requester will wait for a response.
E_D_TOV	The Error_Detect_Timeout Value used for FxPorts as the timeout value for detecting an error condition.
F_S_TOV	The Fabric_Stability_Timeout Value used to ensure that fabric stability has been achieved during fabric configuration.
Network Drop Latency	Network latency of this switch. This is the time interval after which the frames are dropped if they are not delivered in the order they were transmitted. Note that network latency is always greater than switch latency.
Switch Drop Latency	The switch latency of this switch. This is the time interval after which a switch drops the undelivered frames on a link which went down after delivering some frames to the next hop. This way the undelivered frames can be transmitted on a new link if there is one available.
InOrderDelivery	The InOrderDelivery guarantee flag of device. If true, then the InOrder Delivery is guaranteed. If false, it is not guaranteed.

Field	Description
TrunkProtocol	Enables or disables the trunking protocol for the device. The trunking protocol is used for negotiating trunk mode and calculating operational VSANs on an EISL link. It also performs port VSAN consistency checks. On non-trunking ISL links, if the port VSANs are different, the E ports will be isolated. To avoid this isolation, this should be set to disable.

WWN Manager

Field	Description
SwitchWWN	The World-Wide Name of this fabric element. It's a 64-bit identifier and is unique worldwide.
Type 1 WWNs	
Max	Maximum number of NAA Type 1 WWNs that are available for assignment to internal entities.
Available	Number of NAA Type 1 WWNs that are currently available for assignment to internal entities.
Reserved	Number of NAA Type 1 WWNs that are reserved for internal purposes.
Type 2 & 5 WWNs	
Max	Maximum number of total WWNs of types NAA Type 2 and Type 5 WWNs available for assignment to internal entities.
Available	Sum of number of NAA Type 2 and Type 5 WWNs currently available for assignment to the internal entities.
Reserved	Number of total WWNs of types NAA Type 2 and Type 5 WWNs reserved for internal purposes.
Enable Secondary when more WWNs needed	
BaseMacAddress	The first MAC address used for generating World Wide Names (WWNs) when the default range of WWNs generated from supervisor MAC address are exhausted.
MacAddressRange	The number of secondary MAC Addresses starting from and including the wwnmSecondaryBaseMacAddress.

NPV Traffic Map

Field	Description
Switch	Name of the switch
Server Interface	Name of the server interface.
External Interface List	The list of interfaces to which the traffic needs to be mapped to.

NPV Load Balance

Field	Description
Switch	Name of the switch.
Enable	Enable or disable displaying NPV related per server interface information

NPV External Interface Usage

Field	Description
Switch	Name of the switch
Server Interface	Interface on the NPV Device that connects to end devices such as hosts or disks. It is also known as F-port, as it operates in F port mode.
External Interface In Use	Interface on the NPV Device that connects to the NPV Core Switch. It is also known as NP-port as it operates in NP port mode.

NP Link

Field	Description
NPIV (core)	Name of the NPIV core switch.
F port	The F port that is connected to the NPIV core switch
NPV	Name of the NPV switch
Speed	An estimate of the interface's current bandwidth in units of 1,000,000 bits per second. If this object reports a value of 'n' then the speed of the interface is between 'n-500,000' to 'n+499,999'.
Rx Util%	Received traffic Utilization %, total number of octets received on the interface over the speed configured on the interface, including framing characters

Field	Description
Rx Bytes	The total number of octets received on the interface, including framing characters.
Tx Util%	Recetransmittedived traffic Utilization %, total number of octets transmitted out of the interface over the speed configured on the interface, including framing characters.
Tx Bytes	The total number of octets transmitted out of the interface, including framing characters.

FCoE

Config

Field	Description
FC Map	The FCoE Mac Address Prefix used to associate the FCoE Node (ENode).
Default FCF Priority	The default FCoE Initialization Protocol (FIP) priority value advertised by the Fibre Channel Forwarder (FCF) to ENodes.
FKA Adv. Period (sec)	The time interval at which FIP Keep Alive (FKA) messages are transmitted to the MAC address of the ENode.

VSAN-VLAN Mapping



Note

This table applies only to N5k switches running version 4.0(1a) and greater.

Field	Description
VSAN Id	The ID of the VSAN.
VLAN Id	The ID of the VLAN.
Oper State	Shows the operational state of this VLAN-VSAN association entry.

VLAN-VSAN Mapping

Field	Description
VSAN Id	The ID of the VSAN.

Field	Description
VLAN Id	The ID of the VLAN.
Oper State	Shows the operational state of this VLAN-VSAN association entry.

FCoE Statistics

Field	Description
Alignment Errors	The count of frames received on a particular interface that are an integral number of octets in length but do not pass the FCS check.
FCS Errors	The count of frames received on a particular interface that are an integral number of octets in length but do not pass the FCS check. This count does not include frames received with frame-too-long or frame-too-short error.
Single Collision Frames	The count of successfully transmitted frames on a particular interface for which transmission is inhibited by a single collision.
Multiple Collision Frames	The count of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collisions.
SQE Test Errors	The number of times the PLS sublayer generated the SQE TEST ERROR message for a particular interface.
Deferred Transmissions	The count of the number of frames for which the first transmission attempt on a particular interface is delayed because the medium is busy.
Late Collisions	The number of times that a collision is detected on a particular interface later than one slot time into the transmission of a packet.
Excessive Collisions	The count of the number of frames for which transmission on a particular interface fails because of excessive collisions. This counter does not increment when the interface is operating in full-duplex mode.
Internal Mac Transmit Errors	The count of the number of frames for which transmission on a particular interface fails because of an internal MAC sublayer transmit error.
Carrier Sense Errors	The number of times that a carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular interface.

Field	Description
Frame Too Longs	The count of number of frames received on a particular interface that exceed the maximum permitted frame size.
Internal Mac Receive Errors	The count of number of frames for which reception on a particular interface fails because of an internal MAC sublayer receive error.
Symbol Errors	For an interface operating at 100 Mb/s, the number of times there was an invalid data symbol when a valid carrier was present

Ficon

FICON VSANs

Field	Description
VSAN ID	Uniquely identifies a VSAN within a fabric.
Host Can Offline SW	If true, it allows the host to put the system offline.
Host Can Sync Time	If true, the host can set the system time.
Port Control by Host	If true, the host is allowed to alter FICON Director connectivity parameters.
Port Control by SNMP	If true, SNMP manager is allowed to alter FICON director connectivity parameters.
CUP Name	The name of the Control Unit Device.
CUP Enable	Indicates whether the Control Unit Device is enabled.
Domain ID	Specifies the domain ID of the switch.
CodePage	The Code Page used in this VSAN.
Character Set	Character set for the code page used in this VSAN.
Active=Saved	If true, the active to saved mode is enabled. All changes will be saved to NVRAM.
User Alert Mode	If true, FICON management stations will prompt on changes.
Device Allegiance	If CUP is in allegiance state with a channel, it cannot accept any commands from any logical paths. A CUP goes in an allegiance state when it accepts command from a channel and forms 'an allegiance' with it until the successful completion of the channel program, at which point the CUP goes in an 'unlocked' mode.

Field	Description
VSAN Time	The system time in the VSAN. This could be set either by the host or be the default global time in the FICON Director. The default global time is the local time in the FICON Director.
VSAN State	Controls the state of the ports belonging to a VSAN in the context of the FICON functionality.
VSAN Serial Number	The serial number of the FICON director for this VSAN.

FICON VSANs Files

Field	Description
Description	Configuration file description.
CUP Name	The name of the Control Unit Device.
Status	Locked indicates no change allowed. Unlocked indicates change allowed.
LastAccessed	The time this file was last accessed.
UserAlertMode	If true, director user alert mode is enabled.

Global

Field	Description
Default Port Prohibited	Check this option to block the default port.

FICON Port Attributes

Field	Description
TypeNumber	The type number for this FICON Director.
SerialNumber	The sequence number assigned to this FICON Director during manufacturing.
Tag	<p>This is the identifier of the peer port.</p> <ul style="list-style-type: none"> • If the peer port's unit type is channel, then PortId will be the CHPID (Channel Path Identifier) of the channel path that contains this peer port. • If the peer port is controlUnit, then PortId will be 0. • If the peer port is fabric, then PortId will be port address of the interface on the peer switch.

Field	Description
FcId	The fabric Id of the other side port (initiator /target). This will be filled only in the case of Fabric ports.
Status	'valid' - if this information is current. 'old' - if this information is cached. Click Clear Old Attributes to clear the cache.
Name	The FICON port name.
Manufacturer	The name of the company that manufactured this FICON Director.
ModelNumber	The model number for this FICON Director.
PlantOfMfg	The plant code that identifies the plant of manufacture of this FICON Director.
UnitType	The peer type of the port that this port is communicating. ==Channel - host ==Control Unit - disk == Fabric - ISL
Alert	Displays one of the following: <ul style="list-style-type: none"> • bitErrThreshExceeded, • lossOfSignalOrSync, • nosReceived, • primitiveSeqTimeOut, • invalidPrimitiveSeq Click Clear to acknowledge and clear this alert.

FICON Port Configuration

Field	Description
Show Installed Ports Only	If true, only physically available ports will be listed in the table.
ESCON Style	ESCON Style Port Configuration display is the Port Configuration table in DM displaying the ESCON Style Ports. In the table, A represents the available ports and P represents the prohibited ports.
Port/ Prohibit	Enter the FICON address of the port and the prohibited list. (This is an alternative to the table grid.)
Name	The port name of this port.
Block	If true, this port will be isolated.
Prohibit Grid	Click on the grid to add or remove the ability of ports to communicate with each other.

FICON Port Numbers

Field	Description
Module	The number of the module in the chassis.
Reserved Port Numbers (Physical)	The reserved port numbers for the module.
NumPorts	The number of ports reserved for that module.
Module Name	The name of the module.
Reserved Port Numbers (Logical)	Chassis slot port numbers. Reserved port numbers for one chassis slot. There can be up to 64 port numbers reserved for each slot in the chassis.

FICON VSANs Director History

To view the latest FICON information, you must click the Refresh button.

Field	Description
KeyCounter	The key counter.
Ports Address Changed	The list of ports that have configuration change for a value of KeyCounter.

Fabric Binding Actions

Field	Description
VSANId	Specifies the unique identifier for a VSAN within a fabric.
Activate	<ul style="list-style-type: none"> activate - results in the valid fabric bindings on this VSAN/VLAN being activated. force activate - results in forced activation, even if there are errors during activation and the activated fabric bindings will be copied to the active database. deactivate - results in deactivation of currently activated valid fabric bindings (if any), on this VSAN/VLAN. Currently active entries (if any), which would have been present in the active database, will be removed. no-selection -
Enabled	The state of activation on this VSAN/VLAN. If true, then an activation has been attempted as the most recent operation on this VSAN/VLAN. If false, then an activation has not been attempted as the most recent operation on this VSAN/VLAN.

Field	Description
Result	Indicates the outcome of the most recent activation/deactivation.
LastChange	When the valid fabric bindings on this VSAN/VLAN were last activated. If the last activation took place prior to the last re-initialization of the agent, then this value will be N/A.
CopyActToConfig	If enabled, results in the active fabric binding database to be copied on to the configuration database on this VSAN/VLAN. Note that the learned entries are also copied.

Fabric Binding Config Database

Field	Description
VSAN Id	Specifies the unique identifier for a VSAN within a fabric.
Peer WWN (Name)	Specifies the switch WWN of a switch that can be part of the fabric.
DomainId	Specifies an insistent domain ID.

Fabric Binding Active Database

Field	Description
VSAN Id	Specifies the unique identifier for a VSAN within a fabric.
Peer WWN	Specifies the switch WWN of a switch that can be part of the fabric.
DomainId	Specifies the insistent domain ID of the switch represented by the corresponding instance of the WWN of a switch.

Fabric Binding Database Differences

Field	Description
VSAN	From the drop down list, select the number VSANs to be compared.

Field	Description
Compare With	<p>Choose the database for comparison:</p> <ul style="list-style-type: none"> • Active - compares the fabric bind active database with respect to configuration database on this VSAN/VLAN. So, the configuration database will be the reference database and the results of the difference operation will be with respect to the configuration database. • Config - compares the fabric bind configuration database with respect to active database on this VSAN/VLAN. So, the active database will be the reference database and the results of the difference operation will be with respect to the active database.
VSAN Id	Specifies the unique identifier for a VSAN within a fabric.
Peer WWN	Specifies the device WWN of a device that can be part of the fabric.
DomainId	Specifies the insistent domain ID of the switch represented by the corresponding instance of the WWN of a switch.
Reason	Indicates the reason for the difference between the databases being compared, for this entry.

Fabric Binding Violations

Field	Description
VSAN Id	Specifies the unique identifier for a VSAN within a fabric.
Peer WWN	The sWWN (switch WWN) of the device that was denied entry into the fabric on one of the local device's ports.
DomainId	The domain ID of the device that was denied entry into the fabric on one of the local device's ports. A value of zero indicates that the switch WWN of the device was not present in the enforced fabric bindings.
DenialTime	When the denial took place.
DenialCount	The number of times this switch has been denied entry into the fabric on one of the local device's ports.
DenialReason	The reason for which the device was denied entry into the fabric on one of the local device's ports.

Fabric Binding Statistics

Field	Description
AllowedReqs	The number of requests from switches to become part of the fabric that have been allowed on this VSAN/VLAN.
DeniedReqs	The number of requests from switches to become part of the fabric that have been denied on this VSAN/VLAN.
Clear	When set to clear, it results in fabric bind statistic counters being cleared on this VSAN/VLAN.

Fabric Binding EFMD Statistics

Field	Description
TxMergeReqs	The number of EFMD Merge Requests transmitted on this VSAN by the local device.
RxMergeReqs	The number of EFMD Merge Requests received on this VSAN by the local device.
TxMergeAccs	The number of EFMD Merge accepts transmitted on this VSAN by the local device.
RxMergeAccs	The number of EFMD Merge accepts received on this VSAN by the local device.
TxMergeRejs	The number of EFMD Merge rejects transmitted on this VSAN by the local device.
RxMergeRejs	The number of EFMD Merge rejects received on this VSAN by the local device.
TxMergeBusys	The number of EFMD Merge Busys transmitted on this VSAN by the local device.
RxMergeBusys	The number of EFMD Merge Busys received on this VSAN by the local device.
TxMergeErrs	The number of EFMD Merge Errors transmitted on this VSAN by the local device.
RxMergeErrs	The number of EFMD Merge Errors received on this VSAN by the local device.

IP Storage

FCIP Profiles

Field	Description
IP Address	The Internet address for this entity.
Port	A TCP port other than the FCIP well-known port on which the FCIP entity listens for new TCP connection requests.
SACK	Whether the TCP Selective Acknowledgement Option is enabled to allow the receiver end to acknowledge multiple lost frames in a single ACK, enabling faster recovery.
KeepAlive (s)	The TCP keep alive timeout for all links within this entity.
ReTrans MinTimeout (ms)	The TCP minimum retransmit timeout for all the links on this entity.
ReTrans Max	The Maximum number of times that the same item of data will be retransmitted over a TCP connection. If delivery is not acknowledged after this number of retransmissions then the connection is terminated.
Send BufSize (KB)	The aggregate TCP send window for all TCP connections on all Links within this entity. This value is used for Egress Flow Control. When the aggregate of the data queued on all connections within this entity reaches this value, the sender is flow controlled.
Bandwidth Max (Kb)	This is an estimate of the Bandwidth of the network pipe used for the B-D product computation, which lets us derive the TCP receive window to advertise.
Bandwidth Min (Kb)	The minimum available bandwidth for the TCP connections on the Links within this entity.
Est Round Trip Time (us)	This is an estimate of the round trip delay of the network pipe used for the B-D product computation, which lets us derive the TCP receive window to advertise.
PMTU Enable	The path MTU discovery.
PMTU ResetTimeout (sec)	The time interval for which the discovered pathMTU is valid, before MSS reverts back to the negotiated TCP value.
CWM Enable	If true, congestion window monitoring is enabled.
CWM BurstSize (KB)	The maximum burst sent after a TCP sender idle period.
Max Jitter	The maximum delay variation (not due to congestion) that can be experienced by TCP connections on this interface.

FCIP Tunnels

Field	Description
-------	-------------

Interface	This identifies the interface on this FCIP device to which this link pertains.
Attached	The interface on which this FCIP link was initiated.
B Port Enable	If true, the B port mode is enabled on the local FCIP link.
B Port KeepAlive	If true, a message is sent in response to a (Fibre Channel) ELS Echo frame received from the peer. Some B Port implementations use ELS Echo request/response frames as Link Keep Alive.
Remote IP Address	The Internet address for the remote FCIP entity.
Remote TCP Port	The remote TCP port to which the local FCIP entity will connect if and when it initiates a TCP connection setup for this link.
Spc Frames Enable	If true, the TCP active opener initiates FCIP special frames and the TCP passive opener responds to the FCIP special frames. If it is set to false, the FCIP special frames are neither generated nor responded to.
Spc Frames RemoteWWN	The World Wide Name of the remote FC Fabric Entity. If this is a zero length string then this link would accept connections from any remote entity. If a WWN is specified then this link would accept connections from a remote entity with this WWN.
Spc Frames Remote Profile Id	The remote FCIP entity's identifier.

FCIP Tunnels (Advanced)

Field	Description
Interface	The interface on which this FCIP link was initiated.
Timestamp Enable	If true, the timestamp in FCIP header is to be checked.
Timestamp Tolerance	The accepted time difference between the local time and the timestamp value received in the FCIP header. By default this value will be EDTOV/2. EDTOV is the Error_Detect_Timeout Value used for Fibre channel Ports as the timeout value for detecting an error condition.
Number Connections	The maximum number of TCP connections allowed on this link.
Passive	If false, this link endpoint actively tries to connect to the peer. If true, the link endpoint waits for the peer to connect to it.
QoS Control	The value to be set for the ToS field in IP header for the TCP control connection.
QoS Data	The value to be set for the ToS field in IP header for the TCP Data connection.
IP Compression	What algorithm is used, if any.
Write Accelerator	The Write accelerator allows for enhancing SCSI write performance.
Tape Accelerator	If true, the tape accelerator (which allows for enhancing Tape write performance) is enabled.
Tape Accelerator Oper	Write Acceleration is enabled for the FCIP link.
TapeRead Accelerator Oper	Enabled automatically when the Tape Accelerator Oper is active.

Field	Description
FlowCtrlBufSize Tape (KB)	The size of the flow control buffer (64K to 32MB). If set to 0, flow control buffer size is calculated automatically by the switch.
IPSec	Indicates whether the IP Security has been turned on or off on this link.
XRC Emulator	Check to enable XRC Emulator. It is disabled by default.
XRC Emulator Oper	Indicates the operational status of XRC Emulator.

FCIP Tunnels (FICON TA)

Field	Description
Interface	A unique value that identifies the interface on this FCIP device to which this link pertains.
VSAN List Admin	The list of VSANs for which FICON Tape Acceleration is configured.
VSAN List Oper	The list of VSANs for which FICON Tape Acceleration is operationally on.

FCIP Tunnels Statistics

Field	Description
Interface	A unique value that identifies the interface on this FCIP device to which this link pertains.
Rx IPCompRatio	The IP compression ratio for received packets on the FCIP device. The value of this object will be presented as a floating point number with two digits after the decimal point.
Tx IPCompRatio	The IP compression ratio for transmitted packets on the FCIP device. The value of this object will be presented as a floating point number with two digits after the decimal point.

FCIP XRC Statistics

Field	Description
ProfileId	Unique ID of the profile.
Interface	Name of the interface.
RRSAccelerated	The number of read record set IUs accelerated.
RRSForwarded	Number of read record set IUs forwarded.
BusyStatus	Number of instances of busy status received from the control unit.

Field	Description
UnitCheckStatus	Number of instances of unit check status received from the control unit.
cfmFcipLinkExtXRCEStatsSelReset	Number of selective resets processed.
BufferAllocErrors	Number of buffer allocation errors.

iSCSI Connection

Field	Description
LocalAddr	The local Internet Network Address used by this connection.
RemoteAddr	The remote Internet Network Address used by this connection.
CID	The iSCSI Connection ID for this connection.
State	<p>The current state of this connection, from an iSCSI negotiation point of view.</p> <ul style="list-style-type: none"> • login - The transport protocol connection has been established, but a valid iSCSI login response with the final bit set has not been sent or received. • full - A valid iSCSI login response with the final bit set has been sent or received. • logout - A valid iSCSI logout command has been sent or received, but the transport protocol connection has not yet been closed.
MaxRecvDSLen	The maximum data payload size supported for command or data PDUs in use within this connection. Note that the size of reported in bytes even though the negotiation is in 512k blocks.
SendMarker	Indicates whether or not this connection is inserting markers in its outgoing data stream.
HeaderDigest	The iSCSI header digest scheme in use within this connection.
DataDigest	The iSCSI data digest scheme in use within this connection.

iSCSI Initiators

Field	Description
Name or IP Address	A character string that is a globally unique identifier for the node represented by this entry.

Field	Description
VSAN Membership	The list of configured VSANs the node represented by this entry can access.
Dynamic	If true, then the node represented by this entry is automatically discovered.
Initiator Type	Indicates whether the node is a host that participates in iSCSI load-balancing.
Persistent Node WWN	If true, then the same FC address is assigned to the node if it were to be represented again in the FC domain with the same node name. Note that the node FC address is either automatically assigned or manually configured.
SystemAssigned Node WWNN	If true, the FC address is automatically assigned to this node. If false, then the FC address has to be configured manually.
Node WWN	The persistent FC address of the node.
Persistent Port WWN	If true, then the same FC address is assigned to the ports of the node if it were to be represented again in the FC domain with the same node name.
Port WWN	All the FC port addresses associated with this node.
AuthUser	This is the only CHAP user name that the initiator is allowed to log in with.
Target UserName	(Optional) The user name to be used for login. If you do not supply a username, the global user name is used.
Target Password	(Optional) The password to be used for login. If you do not supply a password, the global password is used.
Load Metric	A configured load metric of this iSCSI initiator for the purpose of iSCSI load balancing.
Auto Zone Name	The zone name that is used when the system creates automatic zone for this initiator's specific list of targets.

iSCSI Session Initiators

Field	Description
Name or IP Address	The name or IP address of the initiator port.
Alias	The initiator alias acquired at login.

Module Control

Field	Description
Module Id	ID of the module.
Admin Status	Enables or disables the iSCSI feature for the module.
OperStatus	Shows whether the iSCSI interface is enabled or disabled for the module.

iSCSI Global

Field	Description
AuthMethod	The authentication method.
InitiatorIdleTimeout	The time for which the gateway (representing a FC target) waits from the time of last iSCSI session to a iSCSI initiator went down, before purging the information about that iSCSI initiator.
iSLB ZonesetActivate	Checking this option performs automatic zoning associated with the initiator targets
DynamicInitiator	This field determines how dynamic iSCSI initiators are created. Selecting the iSCSI option (default) creates dynamic iSCSI initiators. If you select iSLB then the an iSLB dynamic initiator is created. Selecting the deny option does not allow dynamic creation of the initiators.
Target UserName	The default user name used for login. If an initiator user name is specified, that user name is used instead.
Target Password	The default password used for login. If an initiator password is specified, that password is used instead.

iSCSI Session Statistics

Field	Description
PDU Command	The count of Command PDUs transferred on this session.
PDU Response	The count of Response PDUs transferred on this session.
Data Tx	The count of data bytes that were transmitted by the local iSCSI node on this session.
Data Rx	The count of data bytes that were received by the local iSCSI node on this session.
Errors Digest	Authentication errors.

Field	Description
Errors CxnTimeout	Connection timeouts.

iSCSI Targets

Field	Description
Dynamically Import FC Targets	Check this option to dynamically import FC targets into the iSCSI domain. A target is not imported if it already exists in the iSCSI domain.
iSCSI Name	The iSCSI name of the node represented by this entry.
Dynamic	Indicates if the node represented by this entry was either automatically discovered or configured manually.
Primary Port WWN	The FC address for this target.
Secondary Port WWN	The optional secondary FC address for this target. This is the FC address used if the primary cannot be reached.
LUN Map iSCSI	The configured default Logical Unit Number of this LU.
LUN Map FC Primary	The Logical Unit Number of the remote LU for the primary port address.
LUN Map FC Secondary	The Logical Unit Number of the remote LU for the secondary port address.
Initiator Access All	If true, then all the initiators can access this target even those which are not in the initiator permit list of this target. If false, then only initiators which are in the permit list are allowed access to this target.
Initiator Access List	Lists all the iSCSI nodes that are permitted to access the node represented by this entry. If AllAllowed is false and the value of List is empty, then no initiators are allowed to access this target.
Advertised Interfaces	Lists all the interfaces on which the target could be advertised.
Trespass Mode	The trespass mode for this node. Every iSCSI target represents one or more port(s) on the FC target. If true, the node instructs the FC node to present all LUN I/O requests to secondary port if the primary port is down.
RevertToPrimaryPort	Indicates if it is required to revert back to primary port if the FC target comes back online.

iSCSI iSLB VRRP

Field	Description
VrId, IpVersion	The virtual router number and the IP version (IPv4, IPv6, or DNS).
Load Balance	Indicates whether load balancing is enabled.

iSCSI Initiator Access

Field	Description
Initiator Name	The iSCSI node name.

Initiator Specific Target

Field	Description
Name	A globally unique identifier for the node.
Port WWN(s) Primary	The Fibre-Channel target's port addresses associated with this iSCSI initiator-specific target.
Port WWN(s) Secondary	The Fibre-Channel target's port addresses associated with this iSCSI initiator-specific target.
LUN Map (Hex) iSCSI	The Fibre-Channel target's port addresses associated with this iSCSI initiator-specific target.
LUN Map (Hex) FC Primary	The Fibre-Channel target's port addresses associated with this iSCSI initiator-specific target.
LUN Map (Hex) FC Secondary	The Fibre-Channel target's port addresses associated with this iSCSI initiator-specific target.
No AutoZone Creation	Indicates if a FibreChannel zone is automatically created for this iSCSI initiator-target and the iSCSI initiator. If true the zone is not automatically created. If false (default) the zone is automatically created.
Trespass Mode	The trespass mode for this node. If true the FC node instance presents all LUN I/O requests to the secondary port (fcSecondaryAddress) if the primary port (fcAddress) is down.
Revert to Primary Port	The revert to primary mode for this node. If true the FC node instance presents all LUN I/O requests to the primary port (fcAddress) when the primary port comes back online.
Primary PWWN VSAN	Indicates the VSAN into which the auto zone is placed for this initiator target. If this object is not set then the VSAN is determined by querying the name server.

Field	Description
Secondary PWWN VSAN	Indicates the VSAN into which the auto zone is placed for this initiator target. If this object is not set then the VSAN is determined by querying the name server.

iSCSI Initiator PWWN

Field	Description
Port WWN	The FC address for this entry.

iSCSI Sessions

Field	Description
Type	Type of iSCSI session: <ul style="list-style-type: none"> • normal - session is a normal iSCSI session • discovery - session is being used only for discovery.
TargetName	If Direction is Outbound, this will contain the name of the remote target.
Vsan ID	The VSAN to which this session belongs to.
ISID	The initiator-defined portion of the iSCSI Session ID.
TSIH	The target-defined identification handle for this session.

iSCSI Sessions Detail

Field	Description
ConnectionNumber	The number of transport protocol connections that currently belong to this session.
ImmediateData	Whether the initiator and target have agreed to support immediate data on this session.
Initial	If true, the initiator must wait for a Ready-To-Transfer before sending to the target. If false, the initiator may send data immediately, within limits set by FirstBurstSize and the expected data transfer length of the request.
MaxOutstanding	The maximum number of outstanding Ready-To-Transfers per task within this session.

Field	Description
First	The maximum length supported for unsolicited data sent within this session.
Max	The maximum number of bytes which can be sent within a single sequence of Data-In or Data-Out PDUs.
Sequence	If false, indicates that iSCSI data PDU sequences may be transferred in any order. If true indicates that data PDU sequences must be transferred using continuously increasing offsets, except during error recovery.
PDU	If false, iSCSI data PDUs within sequences may be in any order. If true indicates that data PDUs within sequences must be at continuously increasing addresses, with no gaps or overlay between PDUs.

IP Services

IP Routes

Field	Description
Routing Enabled	When this check box is enabled, the switch is acting as in IP router.
Destination, Mask, Gateway	The value that identifies the local interface through which the next hop of this route should be reached.
Metric	The primary routing metric for this route.
Interface	The local interface through which the next hop of this route should be reached.
Active	Indicates whether the route is active.

IP Statistics ICMP

Field	Description
InParmProbs	The number of ICMP Parameter Problem messages received.
OutParmProbs	The number of ICMP Parameter Problem messages sent.
InSrcQuenchs	The number of ICMP Source Quench messages received.
InRedirects	The number of ICMP Redirect messages received.

Field	Description
InEchos	The number of ICMP Echo (request) messages received.
InEchoReps	The number of ICMP Echo Reply messages received.
InTimestamps	The number of ICMP Timestamp (request) messages received.
InTimestampReps	The number of ICMP Timestamp Reply messages received.
InAddrMasks	The number of ICMP Address Mask Request messages received.
InAddrMaskReps	The number of ICMP Address Mask Reply messages received.
InDestUnreachs	The number of ICMP Destination Unreachable messages received.
InTimeExcds	The number of ICMP Time Exceeded messages received.
OutSrcQuenchs	The number of ICMP Source Quench messages sent.
OutRedirects	The number of ICMP Redirect messages sent. For a host, this value will always be N/A, since hosts do not send redirects.
OutEchos	The number of ICMP Echo (request) messages sent.
OutEchoReps	The number of ICMP Echo Reply messages sent.
OutTimestamps	The number of ICMP Timestamp (request) messages sent.
OutTimestampReps	The number of ICMP Timestamp Reply messages sent.
OutAddrMasks	The number of ICMP Address Mask Request messages sent.
OutAddrMaskReps	The number of ICMP Address Mask Reply messages sent.
OutDestUnreachs	The number of ICMP Destination Unreachable messages sent.
OutTimeExcds	The number of ICMP Time Exceeded messages sent.

IP Statistics IP

Field	Description
InHdrErrors	The number of input data grams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.

Field	Description
InAddrErrors	The number of input data grams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. For entities which are not IP routers and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
InUnknownProtos	The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
InDiscards	The number of input IP data grams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.
OutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (e.g., for lack of buffer space). Note that this counter would include datagrams counted in ipForwDatagrams if any such frames met this (discretionary) discard criterion.
OutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any frames counted in ipForwDatagrams which meet this 'no-route' criterion. Note that this includes any datagrams which a host cannot route because all of its default routers are down.
FragFails	The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be, e.g., because their Don't Fragment flag was set.
ReasmFails	The number of failures detected by the IP re-assembly algorithm (for whatever reason: timed out, errors, etc). Note that this is not necessarily a count of discarded IP fragments since some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received.
InReceives	The total number of input datagrams received from interfaces, including those received in error.
InDelivers	The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).
OutRequests	The total number of IP datagrams which local IP user- protocols (including ICMP) supplied to IP in requests for transmission. Note that this counter does not include any data grams counted in ipForwDatagrams.
ForwDatagrams	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP routers, this counter will include only those frames which were Source-Routed via this entity, and the Source-Route option processing was successful.
FragOKs	The number of IP datagrams that have been successfully fragmented at this entity.
FragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
ReasmReqds	The number of IP fragments received which needed to be reassembled at this entity.
ReasmOKs	The number of IP datagrams successfully re-assembled.

IP Statistics SNMP

Field	Description
BadVersions	The total number of SNMP messages which were delivered to the SNMP entity and were for an unsupported SNMP version.
BadCommunityNames	The total number of SNMP messages delivered to the SNMP entity which used a SNMP community name not known to said entity.
BadCommunityUses	The total number of SNMP messages delivered to the SNMP entity which represented an SNMP operation which was not allowed by the SNMP community named in the message.
ASNParseErrs	The total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages.
TooBig	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is tooBig.
SilentDrops	The total number of GetRequest-PDUs, GetNextRequest-PDUs, GetBulkRequest-PDUs, SetRequest-PDUs, and InformRequest-PDUs delivered to the SNMP entity which were silently dropped because the size of a reply containing an alternate Response-PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the request.
ProxyDrops	The total number of GetRequest-PDUs, GetNextRequest-PDUs, GetBulkRequest-PDUs, SetRequest-PDUs, and InformRequest-PDUs delivered to the SNMP entity which were silently dropped because the transmission of the (possibly translated) message to a proxy target failed in a manner (other than a time-out) such that no Response-PDU could be returned.
NoSuchNames	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is noSuchName.
BadValues	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is badValue.
ReadOnly	The total number valid SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is readOnly. It should be noted that it is a protocol error to generate an SNMP PDU which contains the value readOnly in the error-status field, as such this is provided as a means of detecting incorrect implementations of the SNMP.
GenErrs	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is genErr.
Pkts	The total number of messages delivered to the SNMP entity from the transport service.
GetRequests	The total number of SNMP Get-Request PDUs which have been accepted and processed by the SNMP protocol entity.
GetNexts	The total number of SNMP Get-Next PDUs which have been accepted and processed by the SNMP protocol entity.
SetRequests	The total number of SNMP Set-Request PDUs which have been accepted and processed by the SNMP protocol entity.

Field	Description
OutTraps	The total number of SNMP Trap PDUs which have been generated by the SNMP protocol entity.
OutGetResponses	The total number of SNMP Get-Response PDUs which have been generated by the SNMP protocol entity.
OutPkts	The total number of SNMP Messages which were passed from the SNMP protocol entity to the transport service.
TotalReqVars	The total number of MIB objects which have been retrieved successfully by the SNMP protocol entity as the result of receiving valid SNMP Get-Request and Get-Next PDUs.
TotalSetVars	The total number of MIB objects which have been altered successfully by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs.

IP Statistics UDP

Field	Description
InErrors	The number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port.
InDatagrams	The total number of UDP datagrams delivered to UDP users.
OutDatagrams	The total number of UDP datagrams sent from this entity.
NoPorts	The total number of received UDP datagrams for which there was no application at the destination port.

mgmt0 Statistics

Field	Description
InErrors	Total number of received errors on the interface.
OutErrors	Total number of transmitted errors on the interface.
InDiscards	Total number of received discards on the interface.
OutDiscards	Total number of transmitted discards on the interface.
TotalRxBytes	Total number of bytes received.
TxBytes	Total number of bytes transmitted.
RxFrames	Total number of frames received.
TxFrames	Total number of frames transmitted.

TCP UDP TCP

Field	Description
State	The state of this TCP connection.

TCP UDP UDP

Field	Description
Port	The local port number for this UDP listener.

VRRP General

Field	Description
IP Address Type, VrId, Interface	The IP address type (IPv4, IPv6, or DNS), the virtual router ID, and the interface.
Admin	The admin state of the virtual router (active or notInService).
Oper	The current state of the virtual router. There are three defined values: <ul style="list-style-type: none"> 'initialize', which indicates that all the virtual router is waiting for a startup event. 'backup', which indicates the virtual router is monitoring the availability of the master router. 'master', which indicates that the virtual router is forwarding frames for IP addresses that are associated with this router.
Priority	Specifies the priority to be used for the virtual router master election process. Higher values imply higher priority. A priority of '0' is sent by the master router to indicate that this router has ceased to participate in VRRP and a backup virtual router should transition to become a new master. A priority of 255 is used for the router that owns the associated IP address(es).
AdvInterval	The time interval, in seconds, between sending advertisement messages. Only the master router sends VRRP advertisements.
PreemptMode	Controls whether a higher priority virtual router will preempt a lower priority master.
UpTime	When this virtual router transitioned out of 'initialized'.
Version	The VRRP version on which this VRRP instance is running.

Field	Description
AcceptMode	Controls whether a virtual router in Master state will accept packets addressed to the address owner's IPv6 address as its own if it is not the IPv6 address owner. If true, the virtual router in Master state will accept. If false, the virtual router in Master state will not accept.

VRRP IP Addresses

Field	Description
Interface, VRRP ID, IP Address	Interface, Virtual Router Redundancy Protocol ID, and associated IP address

VRRP Statistics

Field	Description
IP Address Type, Vrid, Interface	The IP address type (IPv4, IPv6, or DNS), the virtual router ID, and the interface.
LastAdvRx	The total number of VRRP advertisements received by this virtual router.
Protocol Traffic MasterIpAddr	The master router's real (primary) IP address. This is the IP address listed as the source in VRRP advertisement last received by this virtual router.
Protocol Traffic BecomeMaster	The total number of times that this virtual router's state has transitioned to MASTER.
Priority 0 Rx	The total number of VRRP frames received by the virtual router with a priority of '0'.
Priority 0Tx	The total number of VRRP frames sent by the virtual router with a priority of '0'.
AuthErrors InvalidType	The total number of frames received with an unknown authentication type.
Other Errors dvIntervalErrors	The total number of VRRP advertisement frames received for which the advertisement interval is different than the one configured for the local virtual router.
Other Errors IpTtlErrors	The total number of VRRP frames received by the virtual router with IP TTL (Time-To-Live) not equal to 255.
Other Errors InvalidTypePktsRcvd	The number of VRRP frames received by the virtual router with an invalid value in the type field.

Field	Description
Other Errors AddressListErrors	The total number of frames received for which the address list does not match the locally configured list for the virtual router.
OtherErrors PacketLengthErrs	The total number of frames received with a frame length less than the length of the VRRP header.
RefreshRate	The interval of time between refreshes.

CDP General

Field	Description
Enable	Whether the Cisco Discovery Protocol is currently running. Entries in CacheTable are deleted when CDP is disabled.
MessageInterval sec	The interval at which CDP messages are to be generated. The default value is 60 seconds.
HoldTime sec	The time for the receiving device holds CDP message. The default value is 180 seconds.
LastChange	When the cache table was last changed.
Supported DeviceId Format	Indicates the Device-ID format capability of the device.
DeviceId Format	An indication of the format of Device-ID contained in the corresponding instance of the supported device.

CDP Neighbors

Field	Description
Switch	The Internet address for this entity.
Local Interface	A unique value that identifies the interface on this FCIP device to which this link pertains.
DeviceName	The remote device's name. By convention, it is the device's fully qualified domain name.
DeviceID	The device ID string as reported in the most recent CDP message.
DevicePlatform	The version string as reported in the most recent CDP message.
Interface	The port ID string as reported in the most recent CDP message.
IPAddress	The (first) network-layer address of the device's SNMP-agent as reported in the address TLV of the most recently received CDP message.
NativeVLAN	The remote device's interface's native VLAN, as reported in the most recent CDP message. The value 0 indicates no native VLAN field (TLV) was reported in the most recent CDP message.

Field	Description
PrimaryMgmtAddr	Indicates the (first) network layer address at which the device will accept SNMP messages as reported in the most recently received CDP message.
SecondaryMgmtAddr	Indicates the alternate network layer address at which the device will accept SNMP messages as reported in the most recently received CDP message.

iSNS Profiles

Field	Description
Addr	The address of the iSNS server.
Port	The TCP port of the iSNS server.

iSNS Servers

Field	Description
Name	The name of the iSNS Server.
TcpPort	The TCP port used for iSNS messages. If TCP is not supported by this server, the value is 0.
Uptime	The time the server has been active.
ESI Non Response Threshold	The number of ESI messages that will be sent without receiving a response before an entity is de-registered from the iSNS database.
# Entities	The number of entities registered in iSNS on the server.
# Portals	The number of portals registered in iSNS on the server.
# Portal Groups	The number of portal groups registered in iSNS on the server.
# iSCSI Devices	The number of iSCSI Nodes registered in iSNS on the server.

iSNS Entities

Field	Description
Entity ID	The iSNS entity identifier for the entity.
Last Accessed	The time the entity was last accessed.

iSNS Cloud Discovery

Field	Description
AutoDiscovery	Whether automatic cloud discovery is turned on or off.

Field	Description
DiscoveryDelay	Time duration between successive IP cloud discovery runs.
Discovery	The IP network discovery command to be executed. <ul style="list-style-type: none"> all - Run IP network discovery for all the gigabit ethernet interfaces in the fabric. noOp (default) - no operation is performed.
CommandStatus	The status of the license install / uninstall / update operation. <ul style="list-style-type: none"> success - discovery operation completed successfully nProgress - discovery operation is in progress none - no discovery operation is performed NoIpNetworkNameSpecified - ipCloud name not specified invalidNetworkName - ipCloud is not configured NoIPSPortNameSpecified - gigE port ifindex not specified invalidIPSPortName - invalid gigE port interface generalISNSFailure - General ISNS Server Failure

iSNS Clouds

Field	Description
Id	The ID of the IP cloud.
Switch WWN	The WWN of the switch in this table.

iSNS Cloud Interfaces

Field	Description
Name, Switch WWN, Interface, Address	The name, Switch WWN, interface, and address of the cloud.

Monitor Dialog Controls

Field	Description
Line Chart	Opens a new window with a line chart representation of the data.
Area Chart	Opens a new window with an area chart representation of the data.

Field	Description
Bar Chart	Opens a new window with a bar chart representation of the data.
Pie Chart	Opens a new window with a pie chart representation of the data.
Reset Cumulative Counters	Resets the counters to 0 if the Column Data display mode is set to Cumulative.
Export to File	Opens a standard Save dialog box. The data is saved as a .TXT file.
Print	Opens a standard Print dialog box.
Update Frequency	The interval at which the data is updated in the monitor dialog.
Column Data	<p>Specifies the type of data that is displayed in the monitor dialog.</p> <ul style="list-style-type: none"> • Absolute Value - Displays the total amount since the switch was booted. This is the default for error monitoring. • Cumulative - Displays the total amount since the dialog was opened. You can reset the counters by clicking the Reset Cumulative Counters button, to gather a new set of cumulative data. • Minimum/sec - Displays the minimum value per second at every refresh interval. • Maximum/sec - Displays the maximum value per second at every refresh interval. • Last Value/sec - Displays the most recent value per second at every refresh interval. This is the default setting for traffic monitoring.
Elapsed	The amount of time that has elapsed since the dialog was opened. You can reset this counter by clicking the Reset Cumulative Counters button, to gather a new set of cumulative data.

iSNS Details iSCSI Nodes

Field	Description
Name	The iSCSI Name of the initiator or target associated with the storage node.
Type	The Node Type bit-map defining the functions of this iSCSI node, where 31 is a Target, 30 is an Initiator, 29 is a Control, and all others are reserved.
Alias	The Alias name of the iSCSI node.

Field	Description
ScnBitmap	The State Change Notification (SCN) bitmap for a node.
WWN Token	An optional globally unique 64-bit integer value that can be used to represent the World Wide Node Name of the iSCSI device in a Fibre Channel fabric.
AuthMethod	The iSCSI authentication method enabled for this iSCSI Node.

iSNS Details Portals

Field	Description
Addr	The Internet address for this portal.
TcpPort	The port number for this portal.
SymName	The optional Symbolic Name for this portal.
EsiInterval	The Entity Status Inquiry (ESI) Interval for this portal.
TCP ESI	The TCP port number used for ESI monitoring.
TCP Scn	The TCP port used to receive SCN messages from the iSNS server.
SecurityInfo	Security attribute settings for the portal as registered in the Portal Security Bitmap attribute.

Security

Security Roles

Field	Description
Name	Name of the role. Click the Create button to define a new role. Click the Rules button to define the rules for this role.
Description	Text description of the user role.
VSAN Scope Enable	Enables the ability to limit the role to specified VSANs.
VSAN Scope List	Specify a list of VSANs to which the role is allowed access.
Interface Scope Enable	(Nexus 5000 Series only) Enables the ability to limit the role to specified interfaces.
Interface Scope List	(Nexus 5000 Series only) Specify a list of interfaces to which the role is allowed access.

Security Role Rules



Note This table applies only to Nexus 5000 Series switches.

Field	Description
Rule Order	The rules are applied in numerical order.
Permit?	Indicates whether the rule will permit or deny the operation.
Rule Operation	The rule can specify read-only access or read-write access to the operation.
Rule Element Type	The rule can be applied to a command, a feature, feature group or all. Select all to apply the rule to all commands and features.
Rule Element	The rule element specifies the command, feature or feature group to which the rule applies.
Features/Groups	Click the Features/Groups button to open the feature group manager.

Feature Group Manager



Note This table applies only to Nexus 5000 Series switches.

Field	Description
Name	The name of the feature group.
Add	To create a new feature group, enter a new feature group name in the Name field, and click Add .
Add Feature	To add features to feature groups, select one or more feature group names in the Feature Groups panel, select features in the Features panel, and click Add Feature .
Apply	To save changes, click the Apply button

AAA LDAP Servers

Field	Description
IP Address Type	The IP address type (IPv4, IPv6, or DNS).
Name or IP Address	The name or IP address of the AAA server.
AuthPort	The Authentication port of the AAA server.

Field	Description
TimeOut(s)	The time in seconds between retransmissions to the AAA server. This value overrides value set in the timeout set in the Features tab for this server. If this value is zero, then the value set in the Features tab will be used.
Retransmits	The additional number of times the AAA server should be tried by the AAA client before giving up on the server. This value overrides value set in the Features tab. If this value is zero, then the value set in the Features tab will be used.
Idle Time (m)	The time interval in minutes, at which the system periodically tests the AAA Server by sending test packets to the server. The default value of 0 means that the AAA server is not tested periodically.
TestUser	The user name to be used in the test packets sent to the AAA Server, to test if the server responds to the requests.
TestPassword	The password to be used in test packets sent to the AAA Server to test if the server responds to the requests.
RootDN	The root name that is used for authenticating access to LDAP server database.
RootDNPasswordEncrType	Type of encryption that is used for the RootDNPassword password.
RootDNPassword	The RootDN password to use if you want to perform root binding. Anonymous bind will be performed if you do not enter a RootDN password.
SSL Mode	Specifies whether the TLS tunnel needs be setup or not, before binding with the LDAP server.

AAA Server Groups

Field	Description
Name	The name of the server group.
Protocol	The AAA protocol to which this server group belongs to.
ServerIdList	This represents ordered list of AAA Servers which form this Server Group. The order in which servers occur within the value determines the Server priority in that group. The first one will be 'Primary' and the rest are secondary (others). A Server Group can not exist without any members.

Field	Description
DeadTime	The DeadTime setting for AAA Server Group. This indicates the length of time in minutes that the system will mark the server dead when a AAA server does not respond to an authentication request. During the interval of the dead time, any authentication request that comes up would not be sent to that AAA server that was marked as dead. The default value of 0 means that the AAA servers will not be marked dead if they do not respond.

AAA Search Map

Field	Description
BaseDN	Specifies the name of the base entry in the LDAP hierarchy from where the LDAP server begins the search while processing the authorization request.
Filter	Specifies the name of the LDAP filter to be used for searching the user entry in LDAP server database.
Attribute	Specifies the LDAP attribute to be used as user profile private attribute.

AAA Applications

Field	Description
ServerGroupIdList	This represents ordered list of AAA server groups that are configured for this application to perform AAA functions. The order in which server groups occur within the value determines the Server Group priority in the list.
Local	The 'Local' AAA means all the AAA functions are performed using the local AAA service provided in the device. If enabled, is used only after trying all the server groups in the server group list.
Trivial	<p>'Trivial' AAA is used only after trying all the server groups and 'Local' AAA (if configured). Trivial AAA corresponds to one of the following based on the value of corresponding instance of AAAFunction.</p> <ul style="list-style-type: none"> • User name based authentication, if 'AAAFunction' value is 'authentication' • No Authorization check, if 'AAAFunction' value is 'authorization' • No accounting, if 'AAAFunction' value is 'accounting'.

AAA Defaults

Field	Description
KeyEncrType	The encryption type of the server key.
AuthKey	The key used in encrypting the frames passed between the AAA server and the client. This key must match the one configured on the server.
TimeOut	The time in seconds between retransmissions to the AAA server.
Retransmits	The additional number of times the AAA server should be tried by the AAA client before giving up on the server.
DirectReq	Specifies whether you can choose an AAA server for authentication during login. If true, you can specify the remote AAA server for authentication during login. If you specify the login name as username@hostname, then the authentication request is sent to the remote AAA server hostname with the user name as user name. If false, you cannot specify the remote AAA server for authentication during login.
DeadTime (m)	The DeadTime setting for AAA server group. This indicates the length of time in minutes that the system will mark the server dead when a AAA server does not respond to an authentication request. During the interval of the dead time, any authentication request that comes up would not be sent to that AAA server that was marked as dead. The default value of 0 means that the AAA servers will not be marked dead if they do not respond.

AAA General

Field	Description
AuthTypeMSCHAP	Indicates whether the MSCHAP authentication mechanism should be used for authenticating the user through the remote AAA server during login. If true, MSCHAP authentication is used. If false, the default authentication mechanism is used.
AuthTypeMSCHAPv2	Indicates whether the MSCHAPv2 authentication mechanism should be used for authenticating the user through remote AAA Server during login. If true, MSCHAP authentication is used. If false, the default authentication mechanism is used.



Note You are recommended to change one authentication mechanism at a time otherwise there might be an error. For example, if you want to change MSCHAP to MSCHAPv2, please choose MSCHAP and apply, and then choose MSCHAPv2 and apply.

AAA Statistics

Field	Description
Authentication	
Requests	The number of authentication requests sent to this server since it was made active. Retransmissions due to request timeouts are counted as distinct requests.
Timeouts	The number of authentication requests which have timed out since the server was made active.
Unexpected	The number of unexpected authentication responses received from this server since it was made active.
Errors	The number of server ERROR authentication responses received from this server since it was made active.
Incorrect	The number of authentication responses which could not be processed since the server was made active.
ResponseTime	Average response time for authentication requests sent to this server, excluding timeouts, since system re-initialization.
Successes	The number of authentication transactions with this server which succeeded since it was made active. A transaction may include multiple request retransmissions if timeouts occur. A transaction is successful if the server responds with either an authentication pass or fail.
Failures	The number of authentication transactions with this server which failed since it was made active. A transaction may include multiple request retransmissions if timeouts occur. A transaction failure occurs if maximum resends have been met or the server aborts the transaction.
Authorization	
Requests	The number of authorization requests sent to this server since it was made active. Retransmissions due to request timeouts are counted as distinct requests.

Field	Description
Timeouts	The number of authorization requests which have timed out since the server was made active. A timeout results in a retransmission of the request. If the maximum number of attempts has been reached, no further retransmissions will be attempted.
Unexpected	The number of unexpected authorization responses received from this server since it was made active. An example is a delayed response to a request which had already timed out.
Errors	The number of server ERROR authorization responses received from this server since it was made active. These are responses indicating that the server itself has identified an error with its authorization operation.
Incorrect	The number of authorization responses which could not be processed since the server was made active. Reasons include inability to decrypt the response, invalid fields, or the response is not valid based on the request.
ResponseTime	Average response time for authorization requests sent to this server, excluding timeouts, since system re-initialization.
Successes	The number of authorization transactions with this server which succeeded since it was made active. A transaction may include multiple request retransmissions if timeouts occur. A transaction is successful if the server responds with either an authorization pass or fail.
Failures	The number of authorization transactions with this server which failed since it was made active. A transaction may include multiple request retransmissions if timeouts occur. A transaction failure occurs if maximum resends have been met or the server aborts the transaction.
Accounting	
Requests	The number of accounting requests sent to this server since system re-initialization. Retransmissions due to request timeouts are counted as distinct requests.
Timeouts	The number of accounting requests which have timed out since system re-initialization. A timeout results in a retransmission of the request. If the maximum number of attempts has been reached, no further retransmissions are attempted.

Field	Description
Unexpected	The number of unexpected accounting responses received from this server since system re-initialization. An example is a delayed response to a request which had already timed out.
Errors	The number of server ERROR accounting responses received from this server since system re-initialization. These are responses indicating that the server itself has identified an error with its accounting operation.
Incorrect	The number of accounting responses which could not be processed since system re-initialization. Reasons include inability to decrypt the response, invalid fields, or the response is not valid based on the request.
ResponseTime	Average response time for accounting requests sent to this server, since system re-initialization excluding timeouts.
Successes	The number of accounting transactions with this server which succeeded since system re-initialization. A transaction may include multiple request retransmissions if timeouts occur. A transaction is successful if the server responds with either an accounting pass or fail.
Failures	The number of accounting transactions with this server which failed since system re-initialization. A transaction may include multiple request retransmissions if timeouts occur. A transaction failure occurs if maximum resends have been met or the server aborts the transaction.
Statistics	
State	<p>Current state of the server.</p> <ul style="list-style-type: none"> • up - Server responding to requests • dead - Server failed to respond <p>A server is marked dead if it does not respond after maximum retransmissions. A server is marked up again either after a waiting period or if some response is received from it</p>
Duration Current (csec)	The elapsed time the server has been in its current state.
Duration Previous (csec)	This object provides the elapsed time the server was been in its previous state prior to the most recent state. This value is zero if the server has not changed state.
TotalDeadTime	The total elapsed time this server's state has had the value dead since system re-initialization.

Field	Description
DeadCount	The number of times this server's state has transitioned to dead since system re-initialization

iSCSI User

Field	Description
iSCSI User	The name of the iSCSI user.
Password	The password of the iSCSI user.

Common Roles



Note Common Roles is not available in displayFCoE mode (use Security Roles).

Field	Description
Description	Description of the common role.
Enable	This specifies whether the common Role has a VSAN restriction or not.
List	List of VSANs user is restricted to.

SNMP Security Users

Field	Description
Role	The user in Security Model independent format.
Password	Password of the common user. For SNMP, this password is used for both authentication and privacy. For CLI and XML, it is used for authentication only.
Digest	The type of digest authentication protocol which is used.
Encryption	The type of encryption authentication protocol which is used.
ExpiryDate	The date on which this user will expire.
SSH Key File Configured	Specifies whether the user is configured with SSH public key.

Field	Description
SSH Key File	<p>The name of the file storing the SSH public key. The SSH public key is used to authenticate the SSH session for this user. Note that this applies to only CLI user. The format can be one of the following:</p> <ul style="list-style-type: none"> • SSH Public Key in OpenSSH format • SSH Public Key in IETF SECSH (Commercial SSH public key format) • SSH Client Certificate in PEM (privacy-enhanced mail format) from which the public key is extracted • SSH Client Certificate DN (Distinguished Name) for certificate based authentication
Creation Type	The type of the credential store of the user. When a row is created in this table by a user, the user entry is created in a credential store local to the device. In case of remote authentication mechanism like AAA Server based authentication, credentials are stored in other (remote) system/device.
Expiry Date	The date on which this user will expire.

SNMP Security Communities

Field	Description
Community	The community string.
Role	The Security Model name.

Security Users Global

Field	Description
Enforce SNMP Privacy Encryption	Specifies whether the SNMP agent enforces the use of encryption for SNMPv3 messages globally on all the users in the system.
Cache Timeout	This specifies maximum timeout value for caching the user credentials in the local system.



Note The privacy password and authentication password are required for an administrator to create a new user or delete an existing user in Device Manager. However, if the administrator does not provide these credentials at the time of creating a new user, Device Manager uses the authentication password of the administrator as the privacy password. If the privacy protocol defined for the user is not DES (default), the SNMP Agent in the MDS will not be able to decrypt the packet and the SNMP Agent times out. If the privacy protocol defined for the user is not DES, the user needs to provide both the privacy password and the protocol when logging in.

FC-SP General/Password

Field	Description
Timeout	Timeout period for FC-SP messages
HashList	Contains a proposed hash mechanism, in the order of preference. The first is the most preferred and the last contains the least preferred.
GroupList	Each ':' separated token contains a value, corresponding to a Diffie-Hellman group identifier.
GenericPasswd	Password for the switch

FC-SP Interfaces

Field	Description
Mode	<p>The FC-SP mode on this interface.</p> <ul style="list-style-type: none"> • If autoPassive, a port would not initiate any FC-SP authentication exchange; but would always take part in FC-SP authentication exchange initiated by the other side. • If autoActive, a port would always try to initiate FC-SP authentication exchange after ESC. If other side does not support FC-SP authentication, port will still be brought up. • If on, port would always try to initiate FC-SP authentication exchange and authentication is done before the port becomes up. If other side does not support FC-SP authentication, port will not be brought up. • If off, port would never initiate FC-SP authentication exchange and send reject to any FC-SP authentication message started from other end. If this is not 'off', then port has to support at least one FC-SP authentication protocol. <p>Note You need to configure the FC-SP DHCHAP mode individually on each switch to avoid the timeout error from DCNM.</p>
Reauthenticate Interval (hr)	The time (in hours) for which a port has to wait before trying to re-authenticate the other end.
Reauthenticate Start	Re-authenticate the other end, if this is set to enable.
Auth Successes	The number of times the FC-SP authentication succeeded on this interface.
Auth Fails	The number of times the FC-SP authentication failed on this interface.
Auth Bypasses	The number of times the FC-SP authentication was bypassed on this interface.

FC-SP Local Passwords

Field	Description
Local WWN	The World Wide Name of the local host.

Field	Description
Password	Password of the local switch.

FC-SP Remote Passwords

Field	Description
Remote WWN	The World Wide Name of the remote host.
Password	Password of the remote switch.

FC-SP Statistics

Field	Description
Auth Succeeded	The number of times the FC-SP authentication succeeded on this interface.
Auth Failed	The number of times the FC-SP authentication failed on this interface.
Auth ByPassed	The number of times the FC-SP authentication was bypassed on this interface.
EspSpiMismatch	The number of frames received with a mismatched SPI.
EspAuthFailed	The number of frames received that failed ESP authentication check.

FC-SP SA (Security Association)

Field	Description
SPI	Displays the Security Parameter Index value.
Salt	Salt used for encryption.
Key	Key used for encryption and authentication.

FC-SP ESP Interfaces

Field	Description
Interface	Name of the interface.

Field	Description
ESP Mode	Specifies the ESP mode as one of the following: <ul style="list-style-type: none"> • None-ESP is not running on the link. • Gcm- Link needs to be encrypted and authenticated. • Gmac-Link needs to be authenticated
EgressSA	Specifies the egress security association to be used. Valid values are between 256 and 65536.
IngressSA1	Specifies the ingress security association to be used. Valid values are between 256 and 65536.
IngressSA2	Specifies the ingress security association to be used. Valid values are between 256 and 65536.
EspFailureReason	Displays the reason of failure. "None" indicates that no error.

PKI General

Field	Description
Switch	Name of the switch.
CertStoreConfig	The certificate store configuration used by the system for authentication.

PKI RSA Key-Pair

Field	Description
Name	The name or label of a key-pair.
Size	The size of the key. The following modulus sizes are defined: <ul style="list-style-type: none"> • 512-bit, 768-bit, 1024-bit, 1536-bit and 2048-bit. Once created, the size cannot be changed. After a key-pair has been deleted through row deletion, the entry can be created again with another size.
FileName	The name of the file storing the RSA private key. This filename is automatically generated from the key-pair name. It is a unix style '/' separated string representing the absolute path of the file in the file system of the device.

Field	Description
Exportable	The key-pair is exportable through the exportpkcs12 PKI support action. Once created, the exportable flag value cannot be changed. After a key-pair has been deleted through row deletion, the entry can be created again with another value for the exportable flag.

PKI Trust Point

Field	Description
Name	The name or label of a trust point.
KeyPair Name	The name of the associated key-pair from a key-pair table. If a key-pair is not yet associated, the value will be a zero length string.
Revoke CheckMethods	<p>Revocation checking methods list which is an ordered list of certificate revocation checking methods to be employed while verifying peer certificates issued by the CA corresponding to this trust point entry. The value of this object is a ordered list of one or more 1-octet values, where each 1-octet value corresponds to a method in the revocation checking method enumeration:</p> <ul style="list-style-type: none"> • none (1) - No revocation status checking needed; instead consider the certificate as not revoked. • crl (2) - Use CRL for checking the revocation status of certificates. • ocsp (3) - Use OCSP for checking the revocation status of certificates. <p>If none occurs in the list, it should be the last value. The octets after the last value in the ordered list should be zero octets.</p> <p>The order in which the revocation checking methods occur within the value of this object determines the order the revocation checking methods are attempted during the verification of a peer certificate. The default value (after row creation) contains only the revocation checking method crl.</p>
OCSPurl	The contact http url of the external OCSP server for certificate revocation checking using OCSP protocol. The default value (after row creation) is a zero length string.

PKI Trust Point Actions

Field	Description
Name	The name or label of the trust point action.
Command	The PKI support action to be triggered for this trust point entry.
Url	Indicates the file name containing the input or output certificate data needed for the PKI support action being triggered on this entry. The file name should be specified as bootflash:<filename> and it should be available on bootflash or get created on bootflash depending upon the action being triggered.
Password	Indicates the password required to perform the PKI support action being triggered. This password is required to be specified only for certreq, importpkcs12 and exportpkcs12 actions. For security reasons, the value, whenever it is retrieved by the management protocol, is always the zero length string.
Last Command	The PKI support action attempted last. The value attempted to be set for cpkiAction object last. If no action has been triggered for the trust point after its creation, then retrieving the value of this object will return none.
Result	The result of the execution of the last PKI support action.

PKI LDAP

Field	Description
Switch	Name of the switch.
Store Type	The type of remote certificate store.
CRL Timer (hrs)	The time interval based on which the CRL's corresponding to the CA certificates are updated. The CA certificates and the corresponding CRL's are fetched from remote certstore for authentication and are stored in local cache to avoid time delays for subsequent authentication.
Server Group Name	The name of the server group that is used for the remote certstore operations.

PKI Certificate Map

Field	Description
Switch	Name of the switch
Filter Name	The unique name of the mapping filter
Subject Name	The subject name of the CA certificate.
Alternate Name Email	AltNameEmail is another unique field and is a part of the subject name, that is used for authentication.
Alternate Name Universal Principal Name	UPN is another unique field and is a part of the subject name, that is used for authentication.

PKI Certificate Map - Application

Field	Description
Switch	Name of the switch.
Purpose / Issuer Name	The issuer name of the certificate
Map Name 1	The name of the first filtering map that will be applied to the certificate with a given purpose and an issuer name.
Map Name 2	The name of the second filtering map that will be applied to the certificate with a given purpose and an issuer name.

PKI Trust Point Detail

Field	Description
Name	The name or label of the key-pair.
IdCert FileName	The name of the file storing the identity certificate. It is a unix style '/' separated string representing the absolute path of the file in the file system of the device. If there is no identity certificate obtained as yet, the value will be a zero length string.
IdCert SubjName	The subject name of the identity certificate. If there is no certificate or no subject name in the certificate, the value of this object will be a zero length string.
IdCert SerialNum	The serial number of the identity certificate. If there is no certificate, the value of this object will be a zero length string.

Field	Description
IdCert StartDate	The time when the identity certificate starts to be valid, corresponding to the notBefore field in the certificate. If there is no certificate, the value of this object will be a zero length string.
IdCert EndDate	The time when the identity certificate validity ends, corresponding to the notAfter field in the certificate. If there is no certificate, the value of this object will be a zero length string.
IdCert FingerPrint	The MD5 fingerprint of the identity certificate in HEX string format. If there is no certificate, the value of this object will be a zero length string.
IssuerCert FileName	The name of the file storing the issuer certificate. It is a unix style '/' separated string representing the absolute path of the file in the file system of the device. If there is no issuer certificate obtained yet, the value of this object will be a zero length string.
IssuerCert SubjName	The issuer name (subject name in issuer certificate which will be the same as the issuer name in the identity certificate if present). If there is no certificate, the value will be a zero length string.
IssuerCert SerialNum	The serial number of the issuer certificate. If there is no certificate, the value will be a zero length string.
IssuerCert StartDate	The time when the issuer certificate starts to be valid, corresponding to the notBefore field in the certificate. If there is no certificate, the value will be a zero length string.
IssuerCert EndDate	The time when the issuer certificate validity ends, corresponding to the notAfter field on in the certificate. If there is no certificate, the value will be a zero length string.
IssuerCert FingerPrint	The MD5 fingerprint of the issuer's certificate in HEX string format. If there is no certificate, the value of this object will be a zero length string.

IKE Global

Field	Description
RemIdentity	Displays the keep alive interval in seconds used by the IKE entity on the managed device with all the peers for the DOI corresponding to this conceptual row.
Key	Displays the type of keep alives to be used by the IKE entity on the managed device with all the peers for the DOI corresponding to this conceptual row.

IKE Pre-Shared AuthKey

Field	Description
KeepAliveInterval (sec)	The Phase 1 ID identity of the peer for which this pre-shared key is configured on the local entity.
IdentityType	The pre-shared authorization key used in authenticating the peer corresponding to this conceptual row.

IKE Policies

Field	Description
Priority	The priority of this ISAKMP Policy entry. The policy with lower value would take precedence over the policy with higher value in the same DOI.
Encr	The encryption transform specified by this ISAKMP policy specification. The Internet KeyExchange (IKE) tunnels setup using this policy item would use the specified encryption transform to protect the ISAKMP PDUs.
Hash	The hash transform specified by this ISAKMP policy specification. The IKE tunnels setup using this policy item would use the specified hash transform to protect the ISAKMP PDUs.
Auth	The peer authentication method specified by this ISAKMP policy specification. If this policy entity is selected for negotiation with a peer, the local entity would authenticate the peer using the method specified by this object.
DHGroup	Specifies the Oakley group used for Diffie Hellman exchange in the Main Mode. If this policy item is selected to negotiate Main Mode with an IKE peer, the local entity chooses the group specified by this object to perform Diffie Hellman exchange with the peer.
Lifetime (sec)	Specifies the lifetime in seconds of the IKE tunnels generated using this policy specification.

IKE Initiator Version

Field	Description
Address	The address of the remote peer corresponding to this conceptual row. This object cannot be modified while the corresponding value of cicIkeCfgInitiatorStatus is equal to active.

Field	Description
Version	The IKE protocol version used when connecting to a remote peer specified in cicIkeCfgInitiatorPAddr. This object cannot be modified while the corresponding value of cicIkeCfgInitiatorStatus is equal to active.

IKE Tunnels

Field	Description
LocalAddress	The address of the local endpoint for the Phase-1 tunnel.
RemoteAddresss	The address of the remote endpoint of the Phase-1 tunnel.
AuthMethod	The authentication method used in Phase-1 negotiations on the control tunnel corresponding to this conceptual row.
Action	The action to be taken on this tunnel. If clear, then this tunnel is cleared. If re-key, then re-keying is forced on this tunnel. The value none would be returned on doing read of this object.

IPSEC Global

Field	Description
Lifetime (sec)	The default lifetime (in seconds) assigned to an IPSEC tunnel as a global policy (maybe overridden in specific cryptomap definitions).
Lifesize (KB)	The default life size in KBytes assigned to an IPSEC tunnel as a global policy (unless overridden in cryptomap definition).

IPSEC Transform Set

Field	Description
Id	This is the sequence number of the transform set that uniquely identifies the transform set. Distinct transform sets must have distinct sequence numbers.
Protocol	Represents the suite of Phase-2 security protocols of this transform set.
ESP Encryption	Represents the transform used for ESP encryption.

Field	Description
ESP Authentication	Represents the transform used to implement integrity check with ESP protocol.
Mode	Represents the encapsulation mode of the transform set.

IPSEC CryptoMap Set Entry

Field	Description
IpFilter	Specifies an IP protocol filter to be secured using this cryptomap entry. When it has a value of zero-length string, it is not valid/applicable.
TransformSetIdList	The list of cipsXformSetId that are members of this CipsStaticCryptomapEntry. The value of this object is a concatenation of zero or more 4-octet strings, where each 4-octet string contains a 32-bit cipsXformSetId value in network byte order. A zero length string value means this list has no members.
AutoPeer	If true the destination address is taken as the peer address, while creating the tunnel.
Peer Address	The IP address of the peer to which this cryptomap entry is currently connected.
PFS	Identifies whether the tunnels instantiated due to this policy item should use Perfect Forward Secrecy (PFS) and if so, what group of Oakley they should use.
LifeTime	Specifies the lifetime of the IPsec Security Associations (SA) created using this IPsec policy entry.
Lifesize Value	Identifies the life size (maximum traffic in bytes that may be carried) of the IPsec SAs created using this IPsec policy entry. When a Security Association (SA) is created using this IPsec policy entry, its life size takes the value of this object.

IPSEC Interfaces

Field	Description
CryptomapName	The index of the static cryptomap table. The value of the string is the name string assigned by the NMS when defining a cryptomap set.
InterfaceList	Interfaces belong to the cryptomap.

IPSEC Tunnels

Field	Description
Local Address	The IP address of the local endpoint for the IPsec Phase-2 tunnel.
RemoteAddress	The type of the IP address of the remote endpoint for the IPsec Phase-2 tunnel.
ESP Encryption	The encryption algorithm used by the outbound security association of the IPsec Phase-2 tunnel.
ESP Encryption KeySize	The key size in bits of the negotiated key to be used with the algorithm denoted by ceipSecTunOutSaEncryptAlgo. For DES and 3DES the key size is respectively 56 and 168. For AES, this will denote the negotiated key size.
ESP Authentication	The authentication algorithm used by the inbound encapsulation security protocol (ESP) security association of the IPsec Phase-2 tunnel.
LifeSize (KB)	The negotiated life size of the IPSEC Phase-2 tunnel in kilobytes.
LifeTime (sec)	The negotiated lifetime of the IPSEC Phase-2 tunnel in seconds. If the tunnel was setup manually, the value of this MIB element should be 0.
Action	The status of the MIB table row.

IP ACL Profiles

Field	Description
Name	This is the unique IP protocol filter profile identifier.
Type	This object determines the usage type for this filter profile. This usage type cannot be changed after the profile has been created.

IP ACL Interfaces

Field	Description
ProfileName	This is the unique IP protocol filter profile identifier.

IP Filter Profiles

Field	Description
Action	If it is set to deny, all frames matching this filter will be discarded and scanning of the remainder of the filter list will be aborted. If it is set to permit, all frames matching this filter will be allowed for further bridging or routing processing.

Field	Description
Protocol	This filter protocol value matches the Internet Protocol Number in the frames. These IP numbers are defined in the Network Working Group Request for Comments (RFC) documents. Setting this to '-1' will make the filtering match any IP number.
Address	The source IP address to be matched for this filter. A value of 0 causes all source address to match.
Mask	This is the wildcard mask for the SrcAddress bits that must match. 0 bits in the mask indicate the corresponding bits in the SrcAddress must match in order for the matching to be successful, and 1 bits are don't care bits in the matching. A value of 0 causes only IP frames of source address the same as SrcAddress to match.
PortLow	If Protocol is UDP or TCP, this is the inclusive lower bound of the transport-layer source port range that is to be matched, otherwise it is ignored during matching. This value must be equal to or less than the value specified for this entry in SrcPortHigh.
PortHigh	If Protocol is UDP or TCP, this is the inclusive upper bound of the transport-layer source port range that is to be matched, otherwise it is ignored during matching. This value must be equal to or greater than the value specified for this entry in SrcPortLow. If this value is '0', the UDP or TCP port number is ignored during matching.
Address	The destination IP address to be matched for this filter. A value of 0 causes all source address to match.
Mask	This is the wildcard mask for the DestAddress bits that must match. 0 bits in the mask indicate the corresponding bits in the DestAddress must match in order for the matching to be successful, and 1 bits are don't care bits in the matching. A value of 0 causes only IP frames of source address the same as SrcAddress to match.
PortLow	If Protocol is UDP or TCP, this is the inclusive lower bound of the transport-layer destination port range that is to be matched, otherwise it is ignored during matching. This value must be equal to or less than the value specified for this entry in PortHigh.

Field	Description
PortHigh	If Protocol is UDP or TCP, this is the inclusive upper bound of the transport-layer destination port range that is to be matched, otherwise it is ignored during matching. This value must be equal to or greater than the value specified for this entry in DestPortLow. If this value is '0', the UDP or TCP port number is ignored during matching.
Precedence	<p>The IP traffic precedence parameters in each frame are used to guide the selection of the actual service parameters when transmitting a datagram through a particular network. Most network treats high precedence traffic as more important than other traffic. The IP Precedence value ranges from '0' to '7', with '7' the highest precedence and '0' the lowest precedence. The value '-1' means to match frames of any IP precedence. In other words, the IP precedence parameter will not be checked if this value is '-1'. The precedence level are:</p> <ul style="list-style-type: none"> • routine(0) - Routine traffic precedence • priority(1) - Priority traffic precedence • immediate(2) - Immediate traffic precedence • flash(3) - Flash traffic precedence • flashOverride(4) - Flash-override traffic precedence • critical(5) - Critical precedence • internet(6) - Internetwork control traffic precedence • network(7) - Network control traffic precedence.
TOS	The Type of Service (TOS) of the frame. The TOS values ranges from '0' to '15'. The value '-1' matches any TOS value.
ICMPType	This filter specifies the ICMP message type to be matched. Setting this value to '-1' will make the filtering match any ICMP message type.
ICMPCode	This filter specifies the ICMP message code to be matched. Setting this value to '-1' will make the filtering match any ICMP code.
TCPEstablished	This filter if true specifies that for TCP protocol, in an established connection, a match occurs if the TCP datagram has the ACK,FIN,PSH,RST,SYN or URG control bits set. If false, a match will occur for any TCP datagram.
LogEnabled	Specifies whether filtered frames will be logged by the filtering subsystem or not. If true, then all frames will be logged. If false, then no frame will be logged.

SSH/Telnet

Field	Description
Enable SSH/Telnet	Check to enable SSH and/or Telnet.
NumBits	The number of bits provided to generate the key. This determines the length of the key string generated by the SSH.
Key	The SSH key string that is generated.
LastCreationTime	The time of the last creation of the key.
Enable	Enables or disables the Secure Shell (SSH) service on the device.

Port Security Actions

Field	Description
Activation	
Action	<ul style="list-style-type: none"> • activate - results in the valid port bindings on this VSAN/VLAN being activated. • activate (Turn LearningOff) - results in the valid port bindings on this VSAN/VLAN being activated and copied to the active database and will also result in auto learn being turned off on this VSAN/VLAN, once the activation is complete. • force activate - results in forced activation, even if there are errors during activation and the activated port bindings will be copied to the active database. • force activate (Turn Learning Off) - results in forced activation along with turning auto learn off after activation and the activated port bindings will be copied to the active database. • deactivate - results in deactivation of currently activated valid port bindings (if any), on this VSAN/VLAN. Currently active entries (if any), which would have been present in the active database, will be removed. • Activation will not be allowed on a VSAN if auto-learn is enabled on that VSAN
Enabled	The state of activation on this VSAN/VLAN. If true, then an activation has been attempted as the most recent operation on this VSAN/VLAN. If false, then an activation has not been attempted as the most recent operation on this VSAN/VLAN.

Field	Description
Result	Indicates the outcome of the most recent activation/deactivation.
Last Change	When the valid port bindings on this VSAN/VLAN were last activated. If the last activation took place prior to the last re-initialization of the agent, then this value will be N/A.
CopyActiveToConfig	If enabled, results in the active port binding database to be copied on to the configuration database on this VSAN/VLAN. Note that the learned entries are also copied.
AutoLearn	Helps to learn the valid port binding configuration of devices/ports logged into the local device on all its ports and populate the above active database with the same. This mechanism of 'learning' the configuration of devices/ports logged into the local device over a period of time and populating the configuration is a convenience mechanism for users. If enabled on a particular VSAN, all subsequent logins (FLOGIs) on that VSAN will be populated in the enforced port binding database, provided it is not in conflict with existing enforced port bindings on that VSAN. When disabled, the mechanism of learning is stopped. The learned entries will however be in the active database.
Clear AutoLearned	
Action	<ul style="list-style-type: none"> • Clear VSAN results in port bind auto-learned entries being cleared on this VSAN. • Clear Interface(s) results in port bind auto-learned entries being cleared on the interface specified on this VSAN.
Interface	Specifies the interface(s) on which the port bind auto-learned entries need to be cleared.

Port Security Config Database

Field	Description
Interface or fWWN	Represents the address of the port on the local device through which the device specified can FLOGI. <ul style="list-style-type: none"> • If fwwn, then the value is the fabric WWN of a port on the local device. • If intfIndex, then a port on the local device is being represented by its interface. • If wildCard, then it represents a wild-card entry. The wild-card represents any port on the local device.
Type	The mechanism to identify a switch port.
WWN	Represents the logging-in device address

Port Security Active Database

Field	Description
Interface or fWWN	The address of a port on the local device.
Type	The mechanism to identify a switch port.== fwwn - the local switch port is identified by Fabric WWN(fWWN). == intfIndex - the local switch port is identified by ifIndex. == wildCard - wild card (any switch port on local device).
WWN	Represents the logging in device address.
IsLearnt	Indicates if this entry is a learnt entry or not.

Port Security Database Differences

Field	Description
CompareWith	Specifies the database for the comparison. <ul style="list-style-type: none"> • configDb - compares the configuration database with respect to active database on this VSAN/VLAN. So, the active database will be the reference database and the results of the difference operation will be with respect to the active database. • activeDb - compares the active database with respect to configuration database on this VSAN/VLAN. So, the configuration database will be the reference database and the results of the difference operation will be with respect to the configuration database.
VSANId	The ID of the VSAN to compare against.
Interface/fWWN	The address of a port on the local device.
Type	The mechanism to identify a switch port. <ul style="list-style-type: none"> • fwwn - the local switch port is identified by Fabric WWN(fWWN). • intfIndex - the local switch port is identified by ifIndex. • wildCard - wild card (any switch port on local device).
WWN	Represents the logging in device address.
Reason	Indicates the reason for the difference between the databases being compared, for this entry.

Port Security Violations

Field	Description
Interface	The fWWN of the port on the local device where the login was denied.
End Device	The pWWN of the device that was denied FLOGI on one of the local device's ports.
Or Switch	The sWWN of the device (if the device happens to be a switch), that was denied entry on one of the local device's ports.
Time	When the login denial took place.
Count	The number of times this particular pWWN/nWWN or sWWN has been denied login on this particular local interface.

Port Security Statistics

Field	Description
AllowedLogins	The number of FLOGI requests that have been allowed on this VSAN/VLAN.
DeniedLogins	The number of FLOGI requests that have been denied on this VSAN/VLAN.
Clear	When set to clear, it results in port bind statistic counters being cleared on this VSAN/VLAN.

IPsec

Field	Description
Interface, CryptomapName	The binding of cryptomap sets to the interfaces of the managed entity.

Events

Call Home General

Field	Description
Contact	The contact person for this switch, together with information on how to contact this person.
PhoneNumber	The phone number of the contact person. The phone number must start with '+' and contains only numeric characters except for space and '-'. Some valid phone numbers are +44 20 8332 9091 +45 44886556 +81-46-215-4678 +1-650-327-2600.
EmailAddress	The email address of the contact person. Some valid email addresses are raj@helpme.com, bob@service.com, mtom@abc.caview.ca.us.
StreetAddress	The mailing address of this switch.
CustomerId	A string, in whatever format is appropriate, to identify the customer.
ContractId	A string, in whatever format is appropriate, to identify the support contract between the customer and support partner.
SiteId	A location identifier of this device.
DeviceServicePriority	The service priority of the device. This determines how fast the device has to be serviced.
Enable	Enables or disables the CallHome infrastructure on the local device.

Call Home Destinations

Field	Description
ProfileName,ID	The destination profile name and identifier.

Field	Description
Type	Transmission method type.
EmailAddress	The email address associated this destination profile. Some examples are raj@helpme.com, bob@service.com, mtom@abc.caview.ca.us.
Http Url	The HTTP URL associated with this destination profile.

Call Home Email Setup

Field	Description
From	The email address that is to be used in the From field when sending the email using SMTP. Some examples are raj@helpme.com, bob@service.com, mtom@abc.caview.ca.us.
ReplyTo	The email address that is to be used in the Reply-To field when sending the email using SMTP. Some examples are raj@helpme.com, bob@service.com, mtom@abc.caview.ca.us.
IP Address Type	The IP address type (IPv4, IPv6, or DNS).
Name or IP Address	Name or IP address of the SMTP server.
Port	TCP port of the SMTP server.

Call Home Alerts

Field	Description
Action	Test - sends a Call Home message TestWithInventory - sends a message with inventory details.
Status	The status of the last callhome action invocation.
FailureCause	The failure cause for the last callhome test invocation.
LastTimeSent	When the last CallHome alert was sent.
NumberSent	The number of CallHome alerts sent.
Every	Time frame for sending the periodic software inventory Call Home message.
Throttling Enable	If checked, enables the message throttling mechanism implemented on the system, to limit the number of callhome messages for a alert type within a time frame. The maximum is 30 in a 2-hour time frame, and any further messages for that alert type are discarded.
Enable	If checked, enables the sending of periodic software inventory callhome messages on the system.

Call Home HTTP Proxy Server

Field	Description
Master	Name of the switch.
Address Type	The type of the HTTP proxy server as represented by the value in the HTTP proxy server address.
Address	The address of the HTTP proxy server.
Port	The port of the HTTP proxy server.
Enable	Enable or disable the use of HTTP proxyserver configured for sending callhome messages over HTTP.

Call Home SMTP Servers

Field	Description
Address Type, Address	IP address of the SMTP server.
Port	TCP port of the SMTP server.
Priority	Priority value

Call Home User Defined Command

Field	Description
User Defined Command	Used to configure user defined commands for the callhome alert group types.

Delayed Traps

Field	Description
Enable	Enable or disable delay traps.
Delay	Delay interval in minutes (valid values are between 1 to 60)

Call Home Profiles

Field	Description
MsgFormat	XML, full text, or short text.
MaxMsgSize	Maximum message size that can be sent to destination pointed to by this destination profile.

Field	Description
MsgLevel	Threshold level, used for filtering alert messages sent to a destination. Callhome alert message with severity level lower than the configured threshold level would not be sent. The default threshold level is debug (1), which means all the alert messages will be sent.
AlertGroups	The list of configured alert groups for this destination profile.

Event Destinations Addresses

Field	Description
Address/Port	IP Address and Port to send event.
Security Name	The SNMP parameters to be used when generating messages to be sent to this address.
Security Model	Is used when generating SNMP messages using this entry.
Inform Type	<ul style="list-style-type: none"> • Trap - unacknowledged event • Inform - acknowledged event.
Inform Timeout	This expected maximum round trip time for communicating with the address.
RetryCount	The number of retries to be attempted when a response is not received for a generated message.
Status	<ul style="list-style-type: none"> • Active—Port is active. • NotInService—Port is out of service.

Event Destinations Security (Advanced)

Field	Description
MPModel	The Message Processing Model to be used when generating SNMP messages using this entry.
SecurityModel	The Security Model to be used when generating SNMP messages using this entry.
SecurityName	Identifies the Principal on whose behalf SNMP messages will be generated using this entry.
SecurityLevel	The Level of Security to be used when generating SNMP messages using this entry.

Event Filters General

Field	Description
FSPF - Nbr State Changes	Specifies whether or not the local switch should issue notification when the local switch learns of a change in the Neighbor's state (state in the FSPF Neighbor Finite State Machine) on an interface on a VSAN.
Domain Mgr - ReConfig Fabrics	Specifies whether or not the local switch should issue a notification on sending or receiving ReConfigureFabric (RCF) on a VSAN.
Zone Server - Request Rejects	Specifies if the Zone Server should issue a notification on rejects.
Zone Server - Merge Failures	Specifies if the zone server should issue a notification on merge failures.
Zone Server - Merge Successes	Specifies if the zone server should issue a notification on merge successes.
Zone Server - Default Zone Behavior Change	Specifies if the zone server should issue a notification if the propagation policy changes.
Zone Server - Unsupp Mode	Specifies if the zone server should issue a notification on unsupp mode changes
FabricConfigServer - Request Rejects	Specifies if the Fabric Configuration Server should issue a notification on rejects.
RSCN - ILS Request Rejects	Specifies if the RSCN module should generate notifications when a SW_RSCN request is rejected.
RSCN - ILS RxRequest Rejects	Specifies if the RSCN module should generate notifications when a SW_RSCN request is rejected.
RSCN - ELS Request Rejects	Specifies if the RSCN module should generate notifications when a SCR or RSCN request is rejected.
FRU Changes	A false value will prevent Field Replaceable Unit (FRU) notifications from being generated by this system.
SNMP - Community Auth Failure	Indicates whether the SNMP entity is permitted to generate authenticationFailure traps.
VRRP	Indicates whether the VRRP-enabled router will generate SNMP traps for events defined in this MIB.
FDMI	Specifies if the FDMI should generate notifications when a registration request is rejected.
License Manager	Indicates whether the system should generate notifications.
Port/Fabric Security	Specifies if the system should generate notifications when a port/fabric security issue arises.

Field	Description
FCC	Specifies whether the agent should generate notifications.
Name Server	If checked, the Name Server generates a notification when a request is rejected. If false, the notification is not generated.

Event Filters Interfaces

Field	Description
EnableLinkTrap	Indicates whether linkUp/linkDown traps should be generated for this interface.

Event Filters Control

Field	Description
Variable	Represents the notification to be controlled.
Descr	Description about the notification.
Enabled	Check to enable notification of the control. Shows the status of the control.



Note You see the Descr column only on switches that runs Cisco NX-OS release 5.0 or later.

Link Incident History

Field	Description
Host Time	The local time on the host.
Switch Time	The local time on the switch.
Port	The port number for the link incidents.
Interface	The Fibre Channel interface in the specified port.
Link Incident	The type of incident that occurred.

RMON Thresholds Controls

Field	Description
AlarmEnable	If true, the RMON alarm feature is enabled. If the RMON feature is disabled, all the RMON alarm related polling are stopped. Note that this is only intended for temporary disabling of RMON alarm feature to ensure that the CPU usage by RMON alarms is not detrimental. For permanent disabling on this feature, it suggested that all the entries in the alarmTable are removed.
MaxAlarms	The maximum number of entries allowed in the alarmTable.

RMON Thresholds 64bit Alarms

Field	Description
Interval	The interval in seconds over which the data is sampled and compared with the rising and falling thresholds. When setting this variable, care should be taken in the case of deltaValue sampling - the interval should be set short enough that the sampled variable is very unlikely to increase or decrease by more than $2^{31} - 1$ during a single sampling interval.
Variable	The variable to be sampled. Only variables that resolve to an ASN.1 primitive type of INTEGER (INTEGER, Integer32, Counter32, Counter64, Gauge, or TimeTicks) may be sampled.
SampleType	The method of sampling the selected variable and calculating the value to be compared against the thresholds. If the value is absoluteValue, the value of the selected variable will be compared directly with the thresholds at the end of the sampling interval. If the value is deltaValue, the value of the selected variable at the last sample will be subtracted from the current value, and the difference compared with the thresholds.

Field	Description
Value	The value of the statistic during the last sampling period. For example, if the sample type is deltaValue, this value will be the difference between the samples at the beginning and end of the period. If the sample type is absoluteValue, this value will be the sampled value at the end of the period. This is the value that is compared with the rising and falling thresholds. The value during the current sampling period is not made available until the period is completed and will remain available until the next period completes.
StartupAlarm	The alarm that may be sent when this entry is first set to valid.
Rising Threshold	A threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single event will be generated.
Rising EventId	The ID of the eventEntry that is used when a rising threshold is crossed.
Falling Threshold	A threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single event will be generated.
Falling EventId	The ID of the eventEntry that is used when a falling threshold is crossed. The eventEntry identified by a particular value of this index is the same as identified by the same value of eventIndex. If there is no corresponding entry in the eventTable, then no association exists. In particular, if this value is N/A, no associated event will be generated, as N/A is not a valid event index.
FailedAttempts	The number of times the alarm variable was polled (in the active state) and no response was received.
Owner	The ID of the user who configured this entry.

RMON Thresholds 32bit Alarms

Field	Description
Interval	The interval in seconds over which the data is sampled and compared with the rising and falling thresholds. When setting this variable, care should be taken in the case of deltaValue sampling - the interval should be set short enough that the sampled variable is very unlikely to increase or decrease by more than $2^{31} - 1$ during a single sampling interval.
Variable	The variable to be sampled. Only variables that resolve to an ASN.1 primitive type of INTEGER (INTEGER, Integer32, Counter32, Counter64, Gauge, or TimeTicks) may be sampled.
SampleType	The method of sampling the selected variable and calculating the value to be compared against the thresholds.
Value	The value of the statistic during the last sampling period.
StartupAlarm	The alarm that may be sent when this entry is first set to valid.
Rising Threshold	A threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single event will be generated.
Rising EventId	The ID of the eventEntry that is used when a rising threshold is crossed.
Falling Threshold	A threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single event will be generated.
Falling EventId	The ID of the eventEntry that is used when a falling threshold is crossed.
FailedAttempts	The number of times the alarm variable was polled (in the active state) and no response was received.
Owner	The ID of the user who configured this entry.

RMON Thresholds Events

Field	Description
Description	A comment describing this event entry.

Field	Description
Type	The type of notification that the probe will make about this event. In the case of log, an entry is made in the log table for each event. In the case of SNMP-trap, an SNMP trap is sent to one or more management stations.
Community	The community string.
LastTimeSent	When this event entry last generated an event. If this entry has not generated any events, this value will be N/A.
Owner	The entity that configured this entry and is therefore using the resources assigned to it.

RMON Thresholds Log

Field	Description
Time	When this log entry was created.
Description	A description of the event that activated this log entry.

Admin

Copy Configuration

Field	Description
From	Specifies the type of file to copy from.
To	Specifies the type of file to copy to.
ServerAddress	The IP address of the server from (or to) which to copy the configuration file.
FileName	The file name (including the path, if applicable) of the file.
Protocol	The protocol to be used for any copy.
UserName	Remote user name.
UserPassword	Remote user password
CopyState	Specifies the state of this config-copy request. The value of this object is instantiated only after the row has been instantiated. For example, after the CopyEntryRowStatus has been made active.
CopyFailCause	The reason why the config-copy operation failed. This object is instantiated only when the CopyState for this entry is in the failed state.

Flash Files

Field	Description
Name	Flash file name as specified by the user copying in the file.

Field	Description
Size (B)	Size of the file in bytes. Note that this size does not include the size of the file system file header.
Modified	Date and time the file was last modified.

Compact Flash

Field	Description
Device	Name of the device.
Partition	Flash partition name used to refer to a partition.
Size	Size of the partition.

License Features

Field	Description
Missing	Represents the number of missing usage licenses of this feature, when one or more installed license files containing this feature's license, are missing in the local system. Under normal condition, the value is 0.
Installed Type	A combination of demo, permanent, counted, unlicensed, inGracePeriod for that license.
Installed Count	Maximum number of concurrent usages of this license feature. This is the cumulative license usage count for this feature from all the installed license files, containing this feature's license information.
Status	Represents the number of current usages of this licensed feature.
ExpiryDate	Expiry date of the licensed feature.
GracePeriod	Represents the grace period left for this feature, in days/seconds. Grace period is the number of seconds either an unlicensed feature or a feature whose license has expired is allowed to run.
Errors	Errors, if any.
DefaultLicenses	The maximum number of concurrent usages of this license feature that is included by default.

License Manager Keys

Field	Description
LastModified	Represents the time when the license file contents was last modified.
Feature	Specifies the installed license file name.
Version	The version number of the license file.
Type	<ul style="list-style-type: none"> • permanent - Indicates permanent license

Field	Description
Count	<ul style="list-style-type: none">• uncounted - Specified the uncounted license for this feature.• counted - Indicates the maximum number of concurrent uses of this licensed feature.

License Manager Install

Field	Description
HostId	Contains the License hostid of the local system. It is used to identify the local system when requesting license(s) for this system.
URI	Represents the location on the local system, from which the license file will be picked for installation. User should have copied the license file provided by CISCO-CCO, by some other means (for example, through CLI) to this location. For example, the value could be 'bootflash:licfile1'. This MUST be set to a valid value before 'install'. For uninstall operation the value is irrelevant.
Target Filename	Represents either the name with which the license file will be installed, or the name of the license file for uninstall.

Status	<p>The status of the license install/uninstall operation:</p> <ul style="list-style-type: none"> • success (1) - install/uninstall operation completed successfully. InProgress (2) - License install/uninstall operation is in progress. • corruptedLicenseFile (3) - License file content is Invalid/Corrupted. • targetLicenseFileAlreadyExist (4) - Target license file name already exist. • invalidLicenseFileName (5) - License file does not exist. • duplicateLicense (6) - License file is already installed. • licenseInUse (7) - Can't uninstall a license file which is in use. • generalLicensingFailure (8) - General error from license Manager. • none (9) - no install/uninstall operation is performed. • licenseExpiryConflict(10) - License exist with a different expiration date for the feature. • invalidLicenseCount(11) - License count is invalid for the feature. • notThisHost (12) - License host-id in the license file doesn't match. • licenseInGraceMore (13) - Number of licenses in grace period exceeds the number in install license file. • licenseFileNotFound (14) - License file not found, for install / uninstall / update operation. • licenseFileMissing (15) - A previously installed license file is found missing. • licenseFileMissing (15) - A previously installed license file is found missing. • invalidLicenseFileExtension (16) - License file does not have a.lic extension. • invalidURI (17) - Invalid license file URI, specified for install operation. • noDemoLicenseSupport (18) - Demo License Not Supported. • invalidPlatform (19) - Invalid Platform
--------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

License Manager Usage

Field	Description
Name	Represents the name of the application which has checked out the feature.
Application	The application which has checked out the feature.

Port Licensing

Field	Description
Id	Displays the License host ID of the local system. It is used to identify the local system when requesting licenses.
Max	Maximum number of concurrent usages of this license.
Used	Represents the current number of usages of this licensed feature.

Feature Set

Field	Description
Name	The name of the feature set.
OpStatus	The current operating status of the feature.
Action	The action executed against the feature set.
LastCommand	The last action triggered for the feature set.
Result	The result of the last action that was applied to the feature set.

Feature Control

Field	Description
Feature Name	The name of the feature.
Status	The current operating status of the feature.
Action	Enable or disable a feature.
LastCommand	The result of the last action for the feature.
Result	The failure reason description for the failed execution of last action triggered for the feature.

NTP Servers

Field	Description
IP Address Type	The IP address type (IPv4 or IPv6) of the peer.
Name or IP Address	The name or IP address of the peer.

Field	Description
Mode	<p>The association mode of the NTP server, with values coded as follows:</p> <p>Peer - A host operating in this mode sends periodic messages regardless of the reachability state or stratum of its peer. By operating in this mode the host, usually a LAN workstation, announces its willingness to be synchronized by, but not to synchronize the peer.</p> <p>Server - This type of association is ordinarily created upon arrival of a client request message and exists only in order to reply to that request, after which the association is dissolved. By operating in this mode the host, usually a LAN time server, announces its willingness to synchronize, but not to be synchronized by the peer.</p>
Preferred	Specifies whether this peer is the preferred one over the others. By default, NTP chooses the peer with which to synchronize the time on the local system. If true, NTP will choose the corresponding peer to synchronize the time with. If multiple entries are true, NTP will choose the first one to be set.

NTP General

Field	Description
Leap	Two-bit code warning of an impending leap second to be inserted in the NTP timescale.
RootDelay	A signed fixed-point number indicating the total round-trip delay in seconds, to the primary reference source at the root of the synchronization subnet.
RootDispersion	The maximum error in seconds, relative to the primary reference source at the root of the synchronization subnet.

Running Processes

Field	Description
Name	The name associated with this process. If the name is longer than 32 characters, it will be truncated to the first 31 characters, and a `*' will be appended as the last character to imply this is a truncated process name.

Field	Description
MemAllocated (B)	The sum of all the dynamically allocated memory that this process has received from the system. This includes memory that may have been returned.
CPU Time (us)	The amount of CPU time the process has used, in microseconds.

Show Startup/Running Config

Field	Description
Startup	Backs up startup configuration of the switch to another computer with the specified file name.
Running	Backs up running configuration of the switch to another computer with the specified file name.
TCP Timeout	The value (in seconds) to wait for establishing TCP connection before timing out. Valid values are 1 to 120. A timeout results in abortion of the back up action.
FileName	To specify the name of the file where backup details are stored.
Compress File	Check the Compress File check box to compress the backup log file.

Show EPLD Version

Field	Description
Image URI	URI of the image.
Result	Version of the the image specified in the URI.

Copy Flash Files

Field	Description
Direction	Specifies the direction for file transfer.
Protocol	The protocol to be used for copy.
ServerAddress	The server address to be used.
RemoteUserName	Remote user name for protocols FTP, SFTP, and SCP.
RemotePassword	Remote user password used by FTP, SFTP or SCP.

Field	Description
Server File	Server file name, either in Flash or on a server, depending on the type of copy command. Mandatory. For a copy from Flash: File name must be of the form [device>:][:] where is a value obtained from FlashDeviceName, is obtained from FlashPartitionName and is the name of a file in Flash. If you copy files using xFTP protocol, server files may need to be located in a path that is relative to xFTP root path. Note You may need to manually modify the file path if required.
Switch File	Switch file name. For a copy to Flash: File name must be of the form {device>:][:] where is a value obtained from FlashDeviceName, is obtained from FlashPartitionName and is any character string that does not have embedded colon characters.

Show Tech Support

Field	Description
TCP Timeout	The number (in seconds) to wait for the CLI before timing out.
FileName	The name of the file where the show tech support information will be captured.
Compress File	Check this check box to compress the text file into a ZIP file.

Show Image Version

Field	Description
Image URL	The URL of the image.
Result	The version of the image at the specified URL.

Show Onboard Log

Field	Description
Filter Log By	
Module Number	Slot number of the card in the chassis.
Start Date	Specify a start time.
End Date	Specify an end time.

Field	Description
Capture Show Onboard Log Output to File	
TCP Timeout	Specify a time-out interval from the drop-down list.
FileName	Name of the log file.
Compress File	Check the Compress File check box to compress the log file.

Summary View

Field	Description
Description	An alias name for the interface, as specified by a network manager. For Port Channel and FCIP, this field will always show members if they are available. For FCIP, this field will show compress if compressed.
VSAN(s)	VSAN membership.
Mode	Operating mode of the port> (See Legend).
Connected To	<p>Attached port. This could be a host, storage, or switch port.</p> <p>Note Device Manager connects and manages one switch at a time. If the switch with NPV switch connection information is stored in the core switch and the NPV switch is selected to view, the Connected To information will not be displayed.</p>
Speed	Maximum bandwidth in Gbps.
Rx	<p>One of the following:</p> <p>Utilization %</p> <p>Number of Bytes</p> <p>Number of Frames</p> <p>Average Frame Size</p>
Tx	<p>One of the following:</p> <p>Utilization %</p> <p>Number of Bytes</p> <p>Number of Frames</p> <p>Average Frame Size</p>

Field	Description
Errors	Total number of Rx and Tx errors on the interface. Types of Rx errors include CRC errors, fragmented framed, unsupported class frames, runt frames, jabber frames, and giant Frames. Types of Tx errors are generally CRC errors, but these are rare. If the Errors field is not empty, there are probably Rx errors. For a more detailed breakdown of the error count, check the Monitor dialog box for appropriate interface.
Discards	Total number of Rx and Tx discards on the interface. Rx frames discarded are generally due to protocol errors. On rare occasions, a frame is received without any hardware errors, but a filtering rule set for the MAC address discards the frame due to a mismatch. Discarded Tx frames can be timeout frame discards (port is offline or not up), or timeout frames that are not sent back to the supervisor (class F/2 frames). If the Discards field is not empty, it is probably due to timeout frames.
Log	If checked, writes the record into the message log on each poll interval.

RLIR ERL

Field	Description
Vsan ID	VSAN Identifier of the port.
FC ID	Fibre Channel identifier of the subscribing Nx_Port.
Format	The device type for which the Nx_Port receives RLIR ELS."
RegType	<p>The subscriber's registration type.</p> <ul style="list-style-type: none"> ConditionalRx - The Nx_Port will be the recipient of a link incident record only if no other recipients from the ERL on the VSAN is chosen. AlwaysRx - The Nx_Port will be always chosen as the recipient of a link incident records.

Preferred Host

Field	Description
Vsan ID	VSAN Identifier of the port.
PreFcid	Preferred Fibre Channel identifier of the subscribing Nx_Port.

Preferred Path

Field	Description
Interface	Represents an interface on the local device on which the matched or classified frame will be forwarded.
VSAN Id	The VSAN ID of this FC route map. An arbitrary integer value that identifies a route in this FC route map. Preference level, which indicates the metric or cost of the preferred path. The lower the number the higher the preference.
DestinationDomain	
FCID	The FC ID that needs to be matched with a source address in a frame for flow classification.
Description	
Primary ISL	
Secondary ISL	

Edit iSCSI Advertised Interfaces

Field	Description
Num	The number of the iSCSI target.
Interface	The interface over which the target is to be advertised.

DNS General

Field	Description
Enable	Enables or disables DNS configuration.
Domain Name	The name of the domain where the DNS server is enabled.

DNS Servers

Field	Description
IP Address	The IP Address of the DNS server.

Cisco Fabric Services (CFS) Features

Field	Description
Globally Enabled	Check this box to allow CFS on this switch to distribute feature configurations to other switches. Uncheck the box to prevent CFS from distributing the configuration to other switches.
Feature	The name of the CFS-capable feature.
Status	Status of the CFS-capable feature.
Command	The action to be triggered for the feature. Actions include: <ul style="list-style-type: none"> • noop - No operation. • enable - Enable CFS distribution on the switch. • disable - Disable CFS distribution on the switch. • commit - Commit changes made since the session began. • abort - Discard changes made, and close the session. • clear - Discard changes made without closing the session.
Type	The last CFS feature scope type used.
VSAN Id	The ID of the VSAN on which this feature is running.
RegionId	The distribution region ID that this CFS capable feature maps to. This region is required to be defined prior to its usage.
View Config Changes As	Determines whether to view the changes as running or pending. A pending configuration exists until a Commit or Abort action is triggered for that feature. If the value is running then all subsequent configuration retrieval for this feature will be from the running configuration on the local device. If the value is pending then all subsequent configuration retrieval for this feature will be from the pending configuration on the local device.
LastCommand	The last action performed on this feature.
Result	Result of the action performed on the CFS-capable feature.

Field	Description
Scope	<p>The value of this object represents the attributes of a CFS-capable feature as registered with the CFS infrastructure.</p> <ul style="list-style-type: none"> • fcFabric - indicates that the CFS based distribution for a feature spans the entire FC (Fibre Channel) fabric • ipNetwork - indicates that the CFS based distribution for a feature spans the entire IP network • vsanScope - indicates that the CFS based distribution for a feature is done on per VSAN basis and restricted to a specific VSAN in a FC (Fibre Channel) fabric
PendingConfOwnerAddr	The address of the device in the fabric where the pending configuration exists for the feature.
Lock Owner Switch	The address of the device in the fabric where the pending configuration exists for the feature within this scope.
Lock Owner UserName	The name of the device in the fabric where the pending configuration exists for the feature within this scope.
Merge Status	<p>The result of the last fabric merge for this feature within the context of the combination of scope type and scope value in the system. The following are the results:</p> <ul style="list-style-type: none"> • Success—Fabric merge completed successfully. • InProgress—Fabric merge in progress. You may get this status when the local device that is a part of fabric engaged in the process of merging with another fabric. • Failure—Fabric merge failed. • Waiting—Waiting for existing merge to complete while the conflicts are being cleared. You may get this status when the local device that is a part of fabric waiting for any conflicts to be resolved before initiating a new instance of fabric merge. • Other—None of the other values of this enumeration.
Master	Select the CFS Master switch.

Cisco Fabric Services (CFS) IP Multicast

Field	Description
IP Address Type	The IP address type (IPv4, IPv6, or DNS).
Multicast Address Domain	The multicast address domain to which the CFS distribution is restricted. There is a default multicast address for both IPv4 and IPv6 through which the keep-alive messages are sent and received to discover the CFS capable switches over IP. All switches with similar multicast address form one CFS-over-IP fabric. The default multicast address for IPv4 is 239.255.70.83 and range supported is [239.255.0.0 - 239.255.255.255] The default multicast address for IPv6 is ff13::7743:4653 and the supported range is [ff13::0000:0000 - ff13::ffff:ffff]
Action	Specifies the current operating mode employed in CFS for distribution over the corresponding type of Internet address. By setting the value of this object to 'enable', CFS will enable its capability to distribute the application data across the fabric over the corresponding type of Internet address. By setting the value of this object to 'disable', CFS will disable its capability to distribute the data across the fabric over the corresponding type of Internet address.

Cisco Fabric Service (CFS) IP Static Peers

Field	Description
IP Static Peer	Specifies the address of a CFS peer device intended for distribution.
DiscStatus	Specifies a user defined peer device intended for CFS distribution.

Cisco Fabric Services (CFS) Feature by Region

Field	Description
Feature	Identifies the name of a CFS-capable feature within a distribution region.
RegionId	Identifies a CFS distribution region.

Cisco Fabric Services (CFS) All Region

Field	Description
RegionId	Identifies a CFS distribution region.

Cisco Fabric Services (CFS) Owner

Field	Description
Feature, VSAN	The name of the CFS-capable feature, and the VSAN in which the feature is enabled or committed.
Name or IP Address	The name or IP address of the switch on which the feature is enabled or committed.
UserName	The name of the user who enabled or committed the feature.
Type	The last CFS feature scope type used.

Cisco Fabric Services (CFS) Merge

Field	Description
Feature	The name of the CFS-capable feature.
CFS Merge Status Value	The result of the last fabric merge that occurred.

Logs

SysLog (Since Reboot)


Note

To see the latest logs, please close and launch the Log dialog. 'Refresh' option is not available for page by page dialog.

Field	Description
Switch Time	The local time on the switch.
Facility	Name of the facility that generated the message.
Severity	The severity of the message.
Event	The name of the event being logged
VSAN Id	The VSAN on which the event occurred.
Host Time	The local time on the host.
Description	A description of the event being logged.

SysLog (Severe Events)

Field	Description
Switch Time	The local time on the switch.
Facility	Name of the facility that generated the message.
Severity	The severity of the message.
Event	The name of the event being logged
VSAN Id	The VSAN on which the event occurred.
Host Time	The local time on the host.
Description	A description of the event being logged.

Accounting Log


Note

To see the latest logs, please close and launch the Log dialog. 'Refresh' option is not available for page by page dialog.

Field	Description
Switch Time	The local time on the switch.
Action	The action that occurred (start, stop, or update).
Protocol & Source	The protocol and the IP address of the source switch.
User	The name of the user.
Description	A description of the action, if applicable.

Switch Logging

Field	Description
ConsoleEnable	Indicate whether the Syslog messages should be sent to the console.
ConsoleMsgSeverity	Minimum severity of the message that are sent to the Console.
TerminalEnable	Indicate whether the Syslog messages should be sent to the terminals.
TerminalMsgSeverity	Minimum severity of the message that are sent to the terminals.
LinecardEnable	Indicate whether the Syslog messages should be generated at the line cards.
LinecardMsgSeverity	Minimum severity of the message that are sent from linecards.
LogFileMsgSeverity	Minimum severity of the message that are sent to the log file.

Field	Description
SyslogLogFileName	Name of file to which the Syslog messages are logged.

Syslog Severity Levels

Field	Description
Facility	Batch process that generates messages.
Severity	Minimum severity of the message that are generated by this Syslog message facility.

Syslog Servers

Field	Description
IPAddress Type	The IP address type (IPv4, IPv6, or DNS).
Name or IP Address	The address of the Syslog server.
MsgSeverity	Minimum severity of the message that are sent to this Syslog server.
Facility	The facility to be used when sending Syslog messages to this server.

End Devices - Hosts

Field	Description
Host Enclosure	Name of the host enclosure
Name	Name of the VMware
IP Address	IP Address of the VMware
CPU Count	CPU Count of the VMware
Memory Size	Memory Size of the VMware
Status	Current status of the VMware.
OS	OS of the VMware.
Data Store	Name of the VMware datastore.
Last Update Time	Time at which the DCNM-SAN Server last updated the VMware.

Intelligent Features – Summary

Field	Description
Switch	IP address of the switch.
Module	Name of the module.
Name	Name of the switch.
IOA	Display enabled if the IOA feature is enabled. The field will be blank if this feature is disabled.
DMM	Display enabled if the DMM feature is enabled. The field will be blank if this feature is disabled.
SANTap	Display enabled if the SANTap feature is enabled. The field will be blank if this feature is disabled.

Data Mobility Manager – Modules

Field	Description for a Job Row	Description for a Session Row
Name	The name of the job.	This field is blank.
ID	System-assigned unique identifier for the job.	The session number within the job.
Mode	Server mode or storage mode.	This field is blank.
Existing Storage	Alias name of the port on the existing storage.	LUN number on the existing storage.
New Storage	Alias name of the port on the new storage.	LUN number on the new storage.
Status	Status of the job. A created or scheduled job has not yet started. An in-progress job is currently performing the migration. A completed or verified job has finished successfully. A stopped, failed or reset job has finished unsuccessfully.	Status of the session.
Time	Date and time that the job is scheduled to start. This field is blank if the job has not been scheduled. If the job is in progress, this field displays the date and time that the job started.	If the session is in progress, this field displays the estimated duration remaining until the session completes. Otherwise, the field is blank.
SSM1	Switch number and slot of the SSM executing the migration job.	Displays On SSM 1 if the session is executing on SSM 1.

Field	Description for a Job Row	Description for a Session Row
SSM2	Switch number and slot of the SSM executing the migration job.	Displays On SSM 2 if the session is executing on SSM 2.
Type	Online or offline migration.	This field is blank.
Rate	Best effort, slow, medium or fast. You set the rate when you configure the migration job.	This field is blank.

Storage Media Encryption

Members

Field	Description
Cluster	SME cluster name.
State	The operational state of the SME cluster.
Master	Identifies the SME cluster master's IP address.
Members	Identifies the IP address of the switch that is a member of the SME cluster.
IsLocal?	Identifies if the switch is a local or remote member of this cluster.

Interfaces

Field	Description
Cluster	Identifies the cluster to which this SME interface belongs.
Interfaces	Identifies the SME interface.
State	Operational state of this SME interface.

Hosts

Field	Description
Host	Fibre-channel port name (P_WWN) of the host Nx_Port.
Cluster	Identifies the cluster to which this host port belongs.

SSM Features

Summary

Field	Description
Switch	Name of the switch on the intelligent module.
Module	Slot number of the intelligent module.
Name	Name of the intelligent module.
IOA	IOA state of the intelligent module.
DMM	DMM state of the intelligent module.
SANTap	SANTap state of the intelligent module.
SE	SE state of the intelligent module.

FCWA

Field	Description
Flow Id	Represents the flow identifier.
Init WWN	Represents the pWWN of the initiator in the flow.
Init VSAN	The VSAN ID of the initiator on which the flow is configured.
Target WWN	Represents the pWWN of the target in the flow.
TargetVSAN	The VSAN ID of the target on which the flow is configured.
WriteAcc	Specifies if write-acceleration feature is enabled for this flow. If set to true it is enabled. If set to false, it is disabled.
BufCount	It specifies the number of buffers to be used for write-acceleration.
Stats Enable	Specifies if the statistics gathering needs to be enabled for this flow. If set to true, then it is enabled. If it is set to false, then it is disabled.
Stats Clear	Assists in clearing the statistics for this flow.
Init Verification	The verification status of the initiator device corresponding to the SCSI flow.
Init Module	The status of the linecard where the SCSI flow initiator device is located.
Target Verification	The verification status of the target device corresponding to the SCSI flow.
Target Module	The status of the linecard where the SCSI flow target device is located.

SSM

Field	Description
StartPort, EndPort, Feature	A table containing feature related information for interfaces. This table gives a list of interfaces that are assigned to different features. The interfaces supported are of the type Fibre Channel.
PartnerImageURI	A collection of objects related to SSM Feature to interface mapping.

MSM

Field	Description
Switch	Name of the switch on the MSM module.
Module, StartNode, EndNode, Feature	A table containing the feature related information, such as the MSM module number, the node range that are assigned to different features.


Note

The difference between MSM (Multiservice Modules) and SSM (Services Module) is that SSM could enable the features per port range on a card. For MSM you have to enable it on the whole card.

SANTap CVT

Field	Description
Node WWN	Represents the node world wide name of the CVT created on the module.
Port WWN	Represents the port world wide name of the CVT created on the module.
Name	The administratively assigned name for this CVT.

SANTap DVT

Field	Description
VSAN Id, Port WWN	Represents the port world wide name of the created DVT. It will be the same as the port world wide name of the real target for which data is to be replicated.
Interface	Represents the port on the module where the DVT will be created.
Target VSAN Id	Represents the VSAN of the real target for which this DVT is being created.
Name	The administratively assigned name for this DVT.

Field	Description
LUNSize Handling	Indicates whether the DVT should use the real target LUN size for the virtual LUN or the max LUN size supported which is 2TB.
IO Timeout (sec)	Represents the IO timeout value associated with the DVT. This object should be set during the DVT creation time and cannot be modified later.
Target IO Timeout (sec)	Represents the target IO timeout value associated with the DVT.

NASB

Field	Description
Control	Specifies the device type for the LUNs exposed by the TPC target. A value of 1 sets the device type to the default value of disk. A value of 2 sets the device type to storage array controller. Other values are reserved for future changes.
Multiple	Specifies whether the TPC target is operating in a single LUN or multi-LUN mode. A value of 1 sets the default mode which is single LUN. A value of 2 sets multi LUN mode in which the TPC target exposes 10 LUNs.

NASB Target

Field	Description
Module, VSAN Id, Processor Id	The unique ID number associated with the TPC target. This ID number is unique within the VSAN in which the TPC target is configured.
Virtual Target Node WWN	The TPC target's node world wide name.
Virtual Target Port WWN	The TPC target's port world wide name.
State	The current state of the TPC target.
XCOPY Num	The total number of xcopy commands processed by the TPC target since the module on which this target has been configured has been online.
XCOPY MinData (KB)	The smallest amount of data in kilobytes transferred by the TPC target in a single xcopy command since the module on which this target has been configured has been online.

Field	Description
XCOPY MaxData (KB)	The largest amount of data in kilobytes transferred by the TPC target in a single xcopy command since the module on which this target has been configured has been online.
XCOPY Avgthruput (KBps)	The average kilobytes per second throughput of the TPC target in processing the xcopy commands.

Virtual Initiator

Field	Description
Processor Id	The DPP ID.
Control	If false, it's the data path. If true, it's the control path.

DMM Rate

Field	Description
Fast(MBps)	Specifies the migration rate value for the fast attribute for a specific module.
Medium(MBps)	Specifies the migration rate value for the medium attribute for a specific module.
Slow(MBps)	Specifies the migration rate value for the slow attribute for a specific module.

FCWA Config Status

Field	Description
Overall	The configuration status for write-acceleration feature for this flow.
Initiator	The initiator configuration status for write-acceleration feature for this flow.
Target	The target configuration status for write-acceleration feature for this flow.

Statistics Status

Field	Description
Overall	The configuration status for statistics feature for this flow.

Field	Description
Initiator	The initiator configuration status for statistics feature for this flow.
Target	The target configuration status for statistics feature for this flow.

Statistics I/O Traffic

Field	Description
IOs Read	The total number of SCSI read operations on this LUN on this flow.
IOs Write	The total number of SCSI write operations on this LUN on this flow.
Blocks Read	The total number of blocks that have been read on this LUN on this flow.
Blocks Write	The total number of blocks that have been written on this LUN on this flow.
Bytes Rx	The total number of octets received in link-level frames on this LUN on this flow.
Bytes Tx	The total number of octets transmitted in link-level frames on this LUN on this flow.
Frames Rx	The total number of link-level FC frames received on this LUN on this flow.
Frames Tx	The total number of link-level frames transmitted on this LUN on this flow.

Statistics I/O Traffic Details

Field	Description
Timeouts Read	The total number of SCSI read operations that have timed out on this LUN on this flow.
Timeouts Write	The total number of SCSI write operations that have timed out on this LUN on this flow.
MaxBlocks Read	The maximum number of blocks read across all read operations on this LUN on this flow.
MaxBlocks Write	The total number of blocks that have been written on this LUN on this flow.
MaxTime Read	The maximum response time over all read operations on this LUN on this flow.

Field	Description
MaxTime Write	The maximum response time over all write operations on this LUN on this flow.
MinTime Read	The minimum response time over all read operations on this LUN on this flow.
MinTime Write	The minimum response time over all write operations on this LUN on this flow.
Active Read	The number of read operations that are currently active on this LUN on this flow.
Active Write	The number of write operations that are currently active on this LUN on this flow.

Statistics SCSI Commands

Field	Description
TestUnitRdys	The number of test unit ready SCSI commands sent on this LUN on this flow.
RepLuns	The number of report LUN SCSI commands sent on this LUN on this flow.
InquiryS	The number of SCSI inquiry commands sent on this LUN on this flow.
RdCapacityS	The number of read capacity SCSI commands sent on this LUN on this flow.
ModeSenses	The number of mode sense SCSI commands sent on this LUN on this flow.
ReqSenses	The number of request sense SCSI commands sent on LUN on this flow.

Statistics SCSI Errors

Field	Description
BusyStatuses	The number of busy SCSI statuses received on this LUN on this flow.
StatusResvConfs	The number of reservation conflicts SCSI status received on this LUN on this flow.
TskSetFulStatuses	The number of task set full SCSI statuses received on this LUN on this flow.
AcaActiveStatuses	The number of ACA active statuses received on this LUN on this flow.

Statistics SCSI Sense Errors

Field	Description
NotRdyErrs	The number of NOT READY SCSI SENSE key errors received on this LUN on this flow. This indicates that the logical unit being addressed cannot be accessed.
MedErrs	The number of MEDIUM ERROR SCSI SENSE key errors received on this LUN on this flow. This indicates that the command terminated with a non-recovered error condition possibly caused by a flaw in the medium.
HwErrs	The number of HARDWARE ERROR SCSI SENSE key errors received on this LUN on this flow. This indicates that the target detected a non-recoverable hardware failure.
IllReqErrs	The number of ILLEGAL REQUEST SCSI SENSE key errors received on this LUN on this flow.
UnitAttErrs	The number of UNIT ATTENTION SCSI SENSE key errors received on this LUN on this flow.
DatProtErrs	The number of DATA PROTECT SCSI SENSE key errors received on this LUN on this flow.
BlankErrs	The number of BLANK CHECK SCSI SENSE key errors received on this LUN on this flow.
CpAbtErrs	The number of COPY ABORTED SCSI SENSE key errors received on this LUN on this flow.
AbtCmdErrs	The number of ABORTED COMMAND SCSI SENSE key errors received on this LUN on this flow.
VolFlowErrs	The number of VOLUME OVERFLOW SCSI SENSE key errors received on this LUN on this flow.
MiscmpErrs	The number of VOLUME OVERFLOW SCSI SENSE key errors received on this LUN on this flow.

Compact

Field	Description
Device	This is the flash device sequence number to index, used within the table of initialized flash devices. The lowest value should be 1. The highest should be less than or equal to the value of the ciscoFlashDevicesSupported object

Field	Description
Partition	This is the flash partition name used to refer to a partition by the system. This can be any alpha-numeric character string of the form AAAAAAAAAAnn, where A represents an optional alpha character and n a numeric character. Any numeric characters must always form the trailing part of the string. The system will use only the numeric portion to map to a partition index. Flash operations get directed to a device partition based on this name. The system has a concept of a default partition. This would be the first partition in the device. The system directs an operation to the default partition whenever a partition name is not specified. The partition name is therefore mandatory except when the operation is being done on the default partition, or the device has just one partition (is not partitioned).
Size	This is the flash partition size. It should be an integral multiple of ciscoFlashDeviceMinPartitionSize. If there is a single partition, this size will be equal to ciscoFlashDeviceSize.



CHAPTER 2

Configuring Cisco DCNM SAN Server

- [Configuring Cisco DCNM-SAN Server, on page 173](#)

Configuring Cisco DCNM-SAN Server

This chapter describes Cisco DCNM-SAN Server, which is a platform for advanced MDS monitoring, troubleshooting, and configuration capabilities. No additional software needs to be installed. The server capabilities are an integral part of the Cisco DCNM-SAN software.

This chapter contains the following sections:

Information About Cisco DCNM-SAN Server

Install Cisco DCNM-SAN Server on a computer that you want to provide centralized MDS management services and performance monitoring. SNMP operations are used to efficiently collect fabric information. The Cisco DCNM-SAN software, including the server components, requires about 60 MB of hard disk space on your workstation. Cisco DCNM-SAN Server runs on Windows 2000, Windows 2003, Windows XP, Solaris 9 and 10, and Red Hat Enterprise Linux AS Release 5.

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

DCNM-SAN Server Features

Cisco DCNM-SAN Server has the following features:

- **Multiple fabric management**—DCNM-SAN Server monitors multiple physical fabrics under the same user interface. This facilitates managing redundant fabrics. A licensed DCNM-SAN Server maintains up-to-date discovery information on all configured fabrics so device status and interconnections are immediately available when you open the DCNM-SAN Client.
- **Continuous health monitoring**—MDS health is monitored continuously, so any events that occurred since the last time you opened the DCNM-SAN Client are captured.
- **Roaming user profiles**—The licensed DCNM-SAN 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 Cisco DCNM-SAN Client and Cisco DCNM-SAN Server.



Note You will not be able to manage a SAN fabric if the DCNM-SAN Server is going through a IP NAT firewall to access the SAN fabric. All the IP addresses that are discovered in a SAN fabric must be directly reachable by the DCNM-SAN Server.

Licensing Requirements For Cisco DCNM-SAN Server

When you install DCNM-SAN, the basic unlicensed version of Cisco DCNM-SAN 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 Cisco DCNM-SAN 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.

Installing and Configuring Cisco DCNM-SAN Server

Before you begin

Prior to running Cisco DCNM-SAN Server, you should create a special Cisco DCNM-SAN administrative user on each switch in the fabric or on a remote AAA server. Use this user to discover your fabric topology.

Procedure

-
- Step 1** Install Cisco DCNM-SAN Client and Cisco DCNM-SAN Server on your workstation. See the [Installing Cisco DCNM-SAN Server, on page 174](#).
 - Step 2** Log in to DCNM-SAN.
 - Step 3** Set Cisco DCNM-SAN Server to continuously monitor the fabric. See the [Managing a Cisco DCNM-SAN Server Fabric, on page 177](#).
 - Step 4** Repeat Step 2 through Step 3 for each fabric that you want to manage through Cisco DCNM-SAN Server.
 - Step 5** Install DCNM-SAN Web Server. See the [Verifying Performance Manager Collections, on page 177](#).
 - Step 6** Verify Performance Manager is collecting data. See the [Verifying Performance Manager Collections, on page 177](#).
-

Installing Cisco DCNM-SAN Server

When you install DCNM-SAN, the basic version of the Cisco DCNM-SAN Server (unlicensed) is installed with it. After you click the DCNM-SAN icon, a dialog box opens and you can enter the IP address of a computer running the Cisco DCNM-SAN Server component. If you do not see the Cisco DCNM-SAN 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 DCNM-SAN without specifying a valid server, you are prompted to start the Cisco DCNM-SAN Server locally.

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

For switches running Cisco MDS 9000 FabricWare, you must install DCNM-SAN from the CD-ROM included with your switch, or you can download DCNM-SAN from Cisco.com.



Note You can have only one instance of Cisco DCNM-SAN Server running on a computer. If you have a DCNM-SAN Standalone version on your computer, you may need to uninstall it before you install Cisco DCNM-SAN Server.

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

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



Note If you are upgrading from an earlier version to 5.0(1a) or later, that is configured with HTTPS to use your own self-provisioned or a third-party issued SSL certificate, make sure that you set the keystore password and then restart the DCNM-SAN Server. To set the keystore password, run \$INSTALLDIR/dcm/fm/bin/encrypt.bat ssl.

Procedure

- Step 1** Click the **Install Management Software** link.
- Step 2** Choose **Management Software > Cisco DCNM-SAN**.
- Step 3** Click the **Installing DCNM-SAN** link.
- Step 4** Click the **FM Installer** link.
You see the welcome message in the Cisco DCNM-SAN Installer window.
- Step 5** Click the **Custom** radio button, and then click Next to begin installation.
- Step 6** Check the **I accept the terms of the License Agreement** check box, and then click **Next**.
You see the Install Options dialog box.
- Step 7** Click the Cisco DCNM-SAN Server (Licensed) radio button to install the server components for Cisco DCNM-SAN Server.
- Step 8** Click Add server to an existing server federation to add the server to a federation.
Note You may need to add the following line in the pg-hba.conf file under # IPv4 local connections in order to allow remote hosts to connect to PostgreSQL database: host all all 0.0.0.0/0 md5 After adding, save the configuration file, restart the PostgreSQL database before you install the second server node.

Note If you are joining more than three DCNM-SAN Servers in a federation, you need to use an Oracle database with the following settings.

Example:

```
C:\Documents and Settings\Administrator>sqlplus /nolog
SQL*Plus: Release 10.2.0.1.0 - Production on Wed Jan 6 17:19:32 2010
Copyright (c) 1982, 2005, Oracle. All rights reserved.
SQL> connect / as sysdba;
Connected.
SQL> alter system set processes=150 scope=spfile;
System altered.
SQL> alter system set open_cursors=500 scope=spfile;
System altered.

SQL> shutdown immediate;
Database closed.
Database dismounted.
ORACLE instance shut down.
SQL> startup;
ORACLE instance started.
Total System Global Area 805306368 bytes
Fixed Size 1453836 bytes
Variable Size 218714356 bytes
Database Buffers 583008256 bytes
Redo Buffers 2129920 bytes
Database mounted.
Database opened.
SQL> show parameter processes;
Total System Global Area 805306368 bytes
Fixed Size 1453836 bytes
Variable Size 218714356 bytes
Database Buffers 583008256 bytes
Redo Buffers 2129920 bytes
Database mounted.
Database opened.
SQL> show parameter processes;
NAME TYPE VALUE
-----
aq_tm_processes integer 0
db_writer_processes integer 4
gcs_server_processes integer 0
job_queue_processes integer 4
log_archive_max_processes integer 2
processes integer 100
```

Step 9 Select an installation folder on your workstation for Cisco DCNM-SAN. On Windows, the default location is **C:\Program Files\Cisco Systems**.

Step 10 Click **Next**.

You see the Database Options dialog box.

Step 11 Click the radio button for either **Install PostgreSQL** or **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 12 If you select Use existing DB, click the radio button for either PostgreSQL 8.1/8.2 or Oracle10g.

Step 13 Click Next in the Database Options dialog box.

You see the Configuration Options dialog box.

Step 14 Click Install to install Cisco DCNM-SAN Server.

What to do next

If you are evaluating one of these Cisco DCNM-SAN Server features and want to stop the evaluation period for that feature, you can do that using Device Manager.

Data Migration in Cisco DCNM-SAN Server

The database migration should be limited to the existing database. Data collision can occur when you merge the data between the several databases.

When you upgrade a non federation mode database to a federation mode database for the first time, the cluster sequence table is filled with the values larger than the corresponding ones in the sequence table and conforming to the cluster sequence number format for that server ID.

Verifying Performance Manager Collections

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

Managing a Cisco DCNM-SAN Server Fabric

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

Selecting a Fabric to Manage Continuously

Procedure

Step 1 Choose **Server > Admin**.

You see the Control Panel dialog box with the Fabrics tab open.

Note The Fabrics tab is only accessible to network administrators.

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 Choose one of the following Admin options:

- a) **Manage Continuously:** The fabric is automatically managed when Cisco DCNM-SAN Server starts and continues to be managed until this option is changed to Unmanage.
- b) **Manage:** The fabric is managed by Cisco DCNM-SAN Server until there are no instances of DCNM-SAN viewing the fabric.
- c) **Unmanage:** Cisco DCNM-SAN Server stops managing this fabric.

Step 3

Click **Apply**.

What to do next**Note**

If you are collecting data on these fabrics using Performance Manager, you should now configure flows and define the data collections.

Cisco DCNM-SAN Server Properties File

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

**Note**

As of Cisco NX-OS Release 4.1(1b) and later, you can optionally encrypt the password in the **server.properties** and the **AAA.properties** files.

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 DCNM-SAN 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 DCNM-SAN Server through e-mail.

The following server properties are added or changed in the Cisco DCNM-SAN Release 3.x and later.

SNMP Specific

- **snmp.preferTCP:** If this option is set to true, TCP is the default protocol for Cisco DCNM-SAN Server to communicate with switches. By default, this setting is **true**. For those switches that do not have TCP enabled, Cisco DCNM-SAN Server uses UDP. The advantage of this setting is the ability to designate one TCP session for each SNMP user on a switch. It also helps to reduce timeouts and increase scalability.



Note If you set this option to false, the same choice must be set in DCNM-SAN. The default value of `snmp.preferTCP` for DCNM-SAN is true.

Performance Chart

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

EMC Call Home

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

Event Forwarding

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

Deactivation

- **deactivate.confirm=deactivate**: Specific Request for User to type a String for deactivation.



Note In a federated server environment, you should not change Cisco DCNM-SAN Server properties by modifying `server.properties` file. You must modify the server.properties using web client menu Admin > Configure > Preferences.

Modifying Cisco DCNM-SAN Server

Fabric Manager Release 2.1(2) or later allows you to modify certain Cisco DCNM-SAN Server settings without stopping and starting the server.

Adding Cisco DCNM-SAN Server Users

Before you begin

You must be a network administrator before you can manage users.

To add a Cisco DCNM-SAN Server user or to change the password for an existing user using DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Click the **Local FM Users** tab in the Control Panel dialog box. You see a list of DCNM-SAN 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.
 - Step 3** Set the username and password for the new user and then click **Apply**.
-

Removing Cisco DCNM-SAN Server Users

To remove a Cisco DCNM-SAN Server user using DCNM-SAN, follow these steps:

Before you begin

You must be a network administrator before you can manage users.

Procedure

-
- Step 1** Click the **Local FM Users** tab in the Control Panel dialog box. You see a list of DCNM-SAN users.
 - Step 2** Click the username 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 Cisco DCNM-SAN Server Username and Password

You can modify the username or password used to access a fabric from DCNM-SAN Client without restarting Cisco DCNM-SAN Server.

Procedure

-
- Step 1** Choose Server > Admin.
You see the Control Panel dialog box with the Fabrics tab open.
 - Step 2** Set the Name or Password for each fabric that you are monitoring with Cisco DCNM-SAN Server.

- Step 3** Click Apply to save these changes.
-

Changing the DCNM-SAN Server Fabric Discovery Username and Password

Procedure

- Step 1** Click Server > Admin in Cisco DCNM-SAN.
You see the Control Panel dialog box with the Fabrics tab open.
- Step 2** Click the fabrics that have updated user name and password information.
- Step 3** From the Admin listbox, select Unmanage and then click Apply.
- Step 4** Enter the appropriate user name and password and then click Apply.
For more information, see the *Performance Manager Authentication*.
-

Changing the Polling Period and Fabric Rediscovery Time

Cisco DCNM-SAN 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 DCNM-SAN Client without restarting Cisco DCNM-SAN Server.

Procedure

- Step 1** Choose Server > Admin.
You see the Control Panel dialog box with the Fabrics tab open.
- Step 2** For each fabric that you are monitoring with Cisco DCNM-SAN Server, set the Polling Interval to determine how frequently Cisco DCNM-SAN Server polls the fabric elements for status and statistics.
- Step 3** For each fabric that you are monitoring with Cisco DCNM-SAN Server, set the Rediscover Cycles to determine how often Cisco DCNM-SAN Server rediscovers the full fabric.
- Step 4** Click Apply to save these changes.
-

Changing the IP Address of the Cisco DCNM-SAN Server

Procedure

- Step 1** Stop the Cisco DCNM-SAN Server.
- Step 2** Change the following parameter in the \$INSTALLDIR/conf/FMServer.conf.
wrapper.app.parameter.4=127.0.0.1

Step 3 Change the following parameter in the \$INSTALLDIR/conf/server.properties.

server.bindaddrs = 127.0.0.1

Step 4 Enter the following command to assign a new IP address.

Example:

```
run $INSTALLDIR/bin/PLMapping.bat -p newipaddress 0
```

Assume \$INSTALLDIR is the top directory of DCNM-SAN installation. The above command is for single server instance, where 0 is the server ID.

Using Device Aliases or FC Aliases

You can change whether DCNM-SAN uses FC aliases or global device aliases from DCNM-SAN Client without restarting Cisco DCNM-SAN Server.

Procedure

Step 1 Choose Server > Admin.

You see the Control Panel dialog box with the Fabrics tab open.

Step 2 For each fabric that you are monitoring with Cisco DCNM-SAN 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.

Server Federation

Server Federation is a distributed system that includes a collection of intercommunicated servers or computers that is utilized as a single, unified computing resource. With Cisco DCNM-SAN Server Federation, you can communicate with multiple servers together in order to provide scalability and easy manageability of data and programs running within the federation. The core of server federation includes several functional units such as Cisco DCNM-SAN Server, embedded web servers, database and DCNM-SAN Client that accesses the servers.

The Cisco DCNM-SAN Server in the federation uses the same database to store and retrieve data. The database is shared among different servers to share common information. A DCNM-SAN Client or DCNM-SAN Web Client can open fabrics from the Cisco DCNM-SAN Server using the mapping table. A fabric can be moved from one logical server to another. A logical server also can be moved from one physical machine to another machine.

Restrictions

- You cannot upgrade more than one Cisco DCNM-SAN Server in an existing federation. If you choose to do so, you may not be able to migrate the Performance Manager statistics and other information on that server.
- You may require to synchronize the time on all the DCNM-SAN Servers in a federated server environment.

Mapping Fabric ID to Server ID

The IP address of the physical server will be mapped to the server ID during the installation of the Cisco DCNM-SAN Server whenever the IP address of the physical server is changed, you need to map the IP address to the server ID using the PLMapping script provided with the Cisco DCNM-SAN Server. Whenever the you open or discover a fabric, the fabric ID will be mapped to the server ID . You can move a fabric to a different server ID using the control panel.

Procedure

- Step 1** Choose Server > Admin.
You see the Control Panel.
 - Step 2** Select the fabric that you want to move to a different server and then click Move.
You see the Move Fabric dialog box.
 - Step 3** You see the fabrics that you selected in the Fabrics to Move list box. From the Move To Server drop-down list select the server you want to move the fabric to.
 - Step 4** Click Move.
-

Opening the Fabric on a Different Server

Procedure

- Step 1** Choose Server > Admin.
You see the Control Panel.
- Step 2** Click Discover.
You see the Discover New Fabric dialog box.
- Step 3** In the Seed Switch list box, enter the IP Address of the seed switch.
- Step 4** In the User Name field, enter the username.
- Step 5** In the password field, enter the password.
- Step 6** From the Auth-Privacy drop-down list, choose the privacy protocol you want to apply.
- Step 7** To open the selected fabric in a different server, select the server ID from the Server drop-down list.
- Step 8** Click Discover.

Note You may receive an error message when you discover a fabric in a federation while another Cisco DCNM-SAN Server is joining the federation. You can discover the fabric on after the installation or upgradation is complete.

Viewing the Sessions in a Federation

Procedure

- Step 1** Choose Server > Admin.
- Step 2** Click the Connected Clients tab.
- You see the Control Panel.
-

Additional References

- Server Federation is a licensed feature. For more information on Cisco DCNM-SAN Server Licensing, see Cisco MDS 9000 Family NX-OS Licensing Guide.
- For more information on deploying Cisco DCNM-SAN Server in a federation, see Cisco Fabric Manager Server Federation Deployment Guide.



CHAPTER 3

Configuring Authentication in Cisco DCNM-SAN

- [Configuring Authentication in Cisco DCNM-SAN, on page 185](#)

Configuring Authentication in Cisco DCNM-SAN

This chapter describes the interdependent software components in Cisco DCNM-SAN that communicate with the switches, authentication steps and the best practices for setting up your fabric and components for authentication.

This chapter contains the following sections:

Information About Cisco DCNM-SAN Authentication


Cisco DCNM-SAN contains multiple components that interact to manage a fabric.

These components include:

- Cisco DCNM-SAN Client
- Cisco DCNM-SAN Server
- Performance Manager
- Interconnected fabric of Cisco MDS 9000 switches and storage devices
- AAA server (optional)

[Figure 1: Cisco DCNM-SAN Authentication Example, on page 186](#) shows an example configuration for these components.

Figure 1: Cisco DCNM-SAN Authentication Example



Administrators launch Cisco DCNM-SAN Client and select the seed switch that is used to discover the fabric. The user name and password used are passed to Cisco DCNM-SAN Server and used to authenticate to the seed switch. If this user name and password are not a recognized SNMP user name and password, either Cisco DCNM-SAN Client or Cisco DCNM-SAN 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 Cisco DCNM-SAN Client and server.



Note You may encounter a delay in authentication if you use a remote AAA server to authenticate Cisco DCNM-SAN or Device Manager.



Note You must allow CLI sessions to pass through any firewall that exists between Cisco DCNM-SAN Client and Cisco DCNM-SAN Server.



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

Best Practices for Discovering a Fabric

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

**Caution**

If the Cisco DCNM-SAN Server's CPU usage exceeds 50 percent, it is recommended that you switch to a higher CPU-class system.

We recommend that you use these best practices for discovering your network and setting up Performance Manager. This ensures that Cisco DCNM-SAN Server has a complete view of the fabric. Subsequent Cisco DCNM-SAN 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 Cisco DCNM-SAN Server using a network administrator or network operator role so that Cisco DCNM-SAN Server has a view of all the VSANs in the fabric. When a VSAN-limited user launches Cisco DCNM-SAN Client, that user sees only the VSANs they are allowed to manage.

**Note**

Cisco DCNM-SAN 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.

Setting Up Discovery for a Fabric

Procedure

- Step 1** Create a special Cisco DCNM-SAN administrative user name in each switch on your fabric with network administrator or network operator roles. Or, create a special Cisco DCNM-SAN 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 Cisco DCNM-SAN administrative user name are the same on all switches in the fabric and that this role has access to all VSANs.
- Step 3** Launch Cisco DCNM-SAN Client using the Cisco DCNM-SAN administrative user. This step ensures that your fabric discovery includes all VSANs.
- Step 4** Set Cisco DCNM-SAN Server to continuously monitor the fabric.
- Step 5** Repeat Step 4 for each fabric that you want to manage through Cisco DCNM-SAN Server.

Performance Manager Authentication

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

Procedure

- Step 1** Click Server > Admin in Cisco DCNM-SAN.

You see the Control Panel dialog box with the Fabrics tab open.

- Step 2** Click the fabrics that have updated user name and password information.
- Step 3** From the Admin listbox, choose Unmanage and then click Apply.
- Step 4** Enter the appropriate user name and password and then click Apply.
- Step 5** From the Admin listbox, choose Manage and then click Apply.
- Step 6** To rediscover the fabric, click Open tab and check the check box(es) next to the fabric(s) you want to open in the Select column.

From Cisco DCNM Release 11.5(1), If one of the fabrics is unlicensed, a banner text appears with following message:

Unlicensed fabric found, go to “License Assignment” tab to assign a valid/honor license fabric for selection.

To assign a license, click **License Assignments** tab, click **Assign all** to assign a valid/honor license for all unlicensed fabrics. For more details, refer to honor license mode section in *Cisco DCNM SAN Management Configuration Guide*.

- Step 7** Click Open to rediscover the fabric. Cisco DCNM-SAN Server updates its user name and password information.
- Step 8** Repeat Steps 3 through 7 for any fabric that you need to rediscover.
- Step 9** Choose **Performance > Collector > Restart** to restart Performance Manager and use the new user name and password.

Cisco DCNM-SAN Web Client Authentication

Cisco DCNM-SAN Web Server does not communicate directly with any switches in the fabric. Cisco DCNM-SAN 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 Cisco DCNM-SAN Web Server.

Procedure

- Step 1** Launch Cisco DCNM-SAN Web Client.
- Step 2** Choose Admin > Management Users > Remote AAA to update the authentication used by Cisco DCNM-SAN Web Client.
- Step 3** Set the authentication mode attribute to radius.
- Step 4** Set the RADIUS server name, shared secret, authentication method, and ports used for up to three RADIUS servers.
- Step 5** Click Modify to save this information.
- Step 6** Launch Cisco DCNM-SAN Web Client.
- Step 7** Choose Admin > Management Users > Remote AAA to update the authentication used by Cisco DCNM-SAN Web Client.
- Step 8** Set the authentication mode attribute to tacacs.
- Step 9** Set the TACACS+ server name, shared secret, authentication method, and port used for up to three TACACS+ servers.

Step 10 Click Modify to save this information.

Step 11 Cisco DCNM-SAN does not support SecureID because it is not compatible with SNMP authentication. Cisco DCNM-SAN uses the same login credentials for all the switches in a fabric. Since SecureID cannot be used more than once for authentication, Cisco DCNM-SAN will not be able to establish a connection to the second switch using a SecureID.



CHAPTER 4

Configuring Cisco DCNM-SAN Client

- [Configuring Cisco DCNM-SAN Client, on page 191](#)

Configuring Cisco DCNM-SAN Client

This chapter describes about the Cisco DCNM-SAN Client, which is a java-based GUI application that provides access to the Cisco DCNM-SAN applications from a remote workstation.

This chapter contains the following sections:

Information About DCNM-SAN Client

Cisco DCNM-SAN 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 Nexus 5000 Series switches, Cisco MDS 9000 Family and third-party switches, hosts, and storage devices.

In addition to complete configuration and status monitoring capabilities for Cisco MDS 9000 Family switches and Cisco Nexus 5000 Series switches, Cisco DCNM-SAN Client provides Fibre Channel troubleshooting tools. You can use these health and configuration analysis tools on the MDS 9000 Family switch or Cisco Nexus 5000 Series switch to perform Fibre Channel ping and traceroute.

Cisco DCNM-SAN Release 4.1(1b) and later provides multilevel security system by adding a *server admin* role that allows access to limited features. The configuration capabilities of a *server admin* is limited to FlexAttach and relevant data.



Note You must use the same release of Cisco DCNM-SAN Client and Cisco DCNM-SAN Server.

Cisco DCNM-SAN Advanced Mode

Advanced mode is enabled by default and provides the full suite of Cisco DCNM-SAN features, including security, IVR, iSCSI, and FICON. To simplify the user interface, from the list box in the upper right corner of the Cisco DCNM-SAN Client, choose Simple. In simple mode, you can access basic MDS 9000 features such as VSANs, zoning, and configuring interfaces. Advanced mode option is not available for *server admin* role.

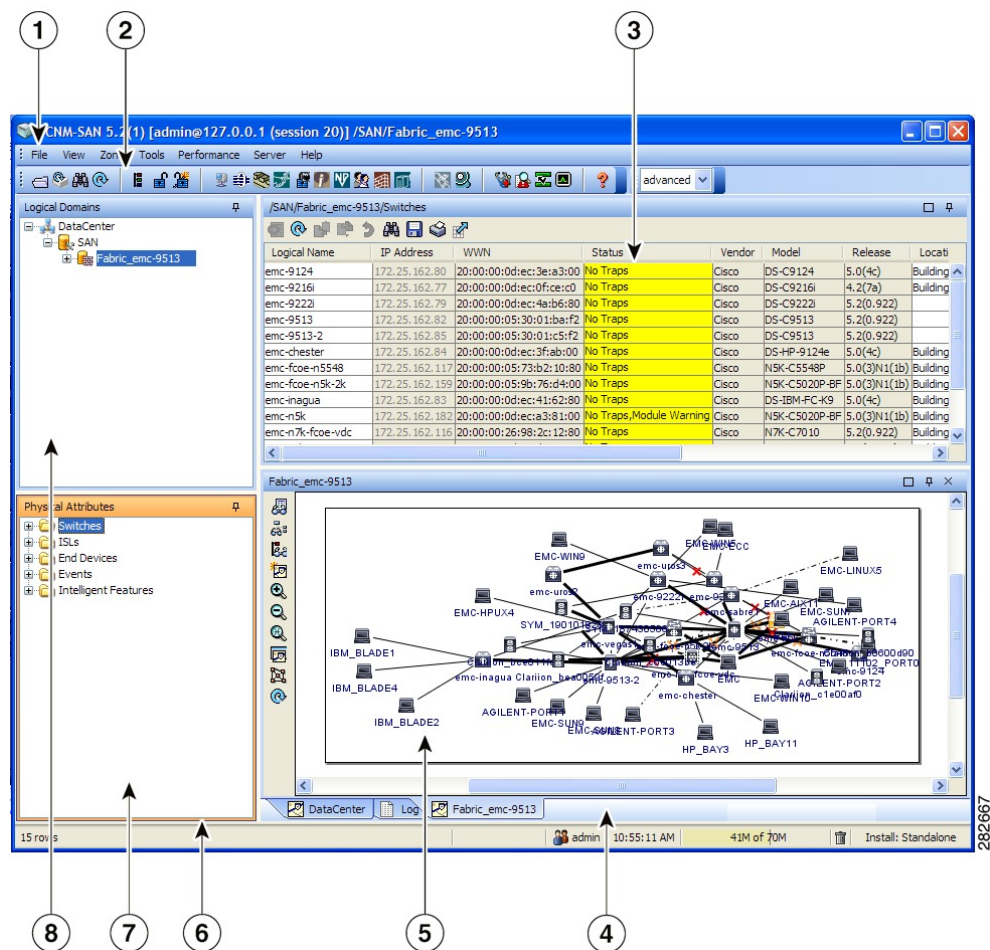
Cisco DCNM-SAN Client Quick Tour: Server Admin Perspective

Cisco DCNM-SAN provides a multilevel security system by adding a *server admin* role that allows access only to limited features. The configuration capabilities of a *server admin* role is limited to FlexAttach and relevant data. The *server admin* can pre-configure SAN for new servers, move a server to another port on the same NPV device or another NPV device and replace a failed server onto the same port without involving the SAN administrator. The *server role admin* will not be able to manage Cisco DCNM-SAN users or connected clients.

Cisco DCNM-SAN provides a much improved user interface by including movable and dockable panes to let users arrange the Physical Attributes pane, Logical Domains pane, Fabric pane and Information pane according to requirements, making it easier to manage the workflow. The dockable panes are also called as dockable frames. A dockable frame can be standalone (floating), minimized or maximized. The logical, physical, information and the fabric panes can be collapsed and expanded as needed. These panes can also be docked at either the right side left side or to the bottom of the workspace.

Cisco DCNM-SAN Main Window

This section describes the Cisco DCNM-SAN Client interface that is specific to *server admin* users as shown in [Figure 4-1](#).



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

Menu Bar



The menu bar at the top of the Cisco DCNM-SAN main window provides options for managing and for controlling the display of information on the Fabric pane. *Server admin* will not have all the options that are available for *SAN admin*. The menu bar provides the following menus:



- File—Opens a new fabric, rediscovers the current fabric, locates switches, sets preferences, prints the map.
- View—Changes the appearance of the map (these options are duplicated on the Fabric pane toolbar).
- Tools—Manages the Server and configuration using the FlexAttach virtual pWWN feature.
- Help—Displays online help topics for specific dialog boxes in the Information pane.

Tool Bar

The Cisco DCNM-SAN main toolbar (specific to *server admin*) provides icons for accessing the most commonly used menu bar options as shown in [Table 1: Cisco DCNM-SAN Client Main Toolbar, on page 193](#).

Table 1: Cisco DCNM-SAN Client Main Toolbar

Icon	Description
	Opens switch fabric.
	Rediscovers current fabric.

Icon	Description
	Finds in the map.
	Shows online help.

Logical Domains Pane

Use the Logical Domains pane to view fabrics and to access user-defined groups. You can expand the groups to see different user-defined groups. The non-editable groups created for each core switch contains their NPV switches.

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 or group.






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





- Switches—Views and configures hardware, system, licensing, and configuration files.
- Interfaces—Views and configures FC physical, FC logical, VFC (FCoE), Ethernet, SVC, and PortChannel interfaces.

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 2: Information Pane Toolbar

Icons	Description
	Applies configuration changes.
	Refreshes table values.
	Copies data from one row to another.
	Pastes the data from one row to another.
	Undoes the most recent change.

Icons	Description
	Finds a specified string in the table.
	Exports and saves information to a file.
	Prints the contents of the Information pane.
	Displays a non-editable copy of the table in the Information pane in its own window, which you can move around the screen.

Fabric Pane

Use the Fabric pane to display the graphical representation of your fabric. Table 5-1 explains the graphics you may see displayed, depending on which devices you have in your fabric.

The bottom of the Fabric pane has the following tabs:

- Fabric—When displaying multiple fabrics, each fabric has its own tab. You can switch between fabrics by clicking on their respective tabs.
- Log—Displays messages that describe Cisco DCNM-SAN 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, SNMP, and FICON.

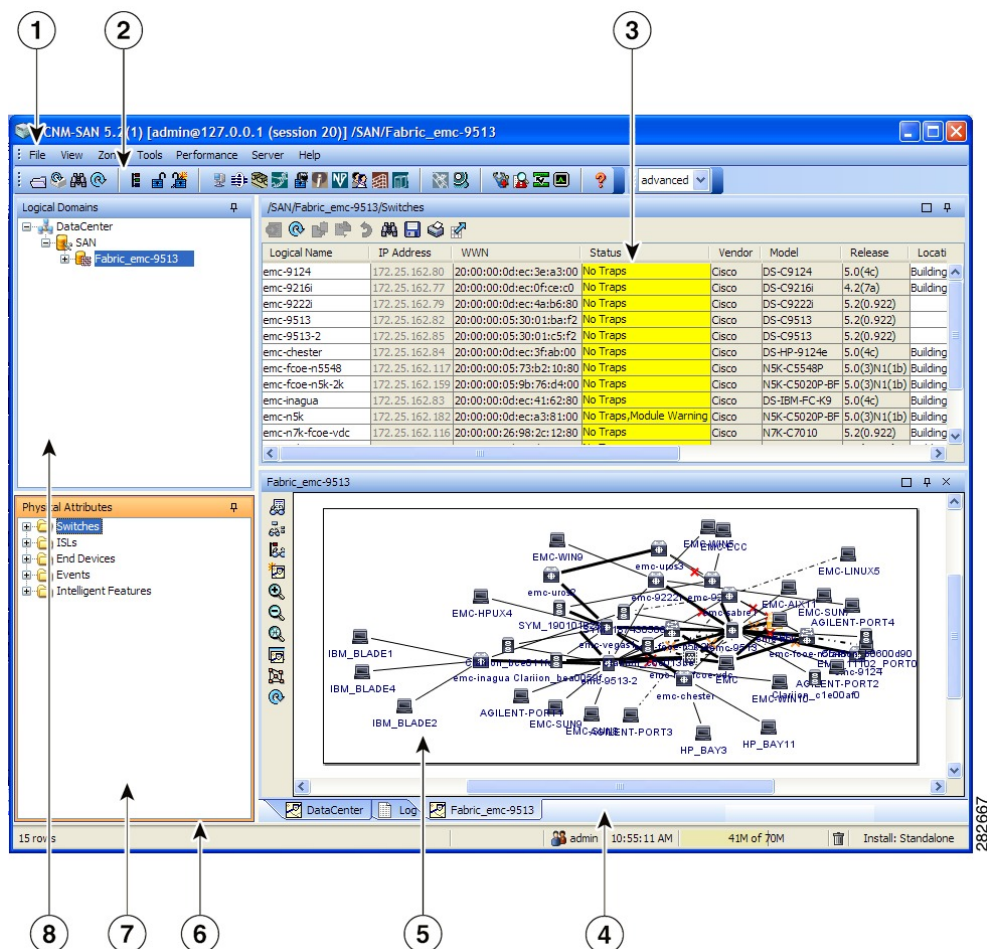


Note Fabric map display is based on what you select in the logical domain pane. When you select a fabric node, all the switches that belong to that fabric will be enabled. When you select the group node, all the switches that belong to the groups listed under that group node will be enabled. When you select only a group, all the switches that belong to the specific group will be enabled.

Cisco DCNM-SAN Client Quick Tour: Admin Perspective

This section describes the Cisco DCNM-SAN Client interface shown in .

Figure 2: Cisco DCNM-SAN Main Window



- 1 Menu bar—Provides access to options that are organized by menus.
- 2 Toolbar—Provides icons for direct access to the most commonly used options on the File, Tools, and Help menus.
- 3 Information pane—Displays information about whatever option is selected in the menu tree.
- 4 Status Bar (right side)—Shows the last entry displayed by the discovery process and the possible error message.
- 5 Fabric pane—Displays a map of the network fabric, including switches, hosts, and storage. It also provides tabs for displaying log and event data.
- 6 Status Bar (left side)—Shows short-term transient messages, such as the number of rows displayed in a table.
- 7 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.

- | | |
|---|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 8 | Logical Domains pane—Displays a tree of configured SAN, fabrics, VSANs, and zones, and provides access to user-defined groups. The label next to the segmented VSAN indicates the number of segments. |
|---|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|



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 Cisco DCNM-SAN 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 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 [Cisco DCNM-SAN Troubleshooting Tools, on page 220](#).
- **Performance**—Runs and configures Performance Manager and Cisco Traffic Analyzer, and generates reports.
- **Server**—Runs administrative tasks on clients and fabrics. Provides Cisco DCNM-SAN Server management and a **purge** command. Lists fabrics being managed.
- **Help**—Displays online help topics for specific dialog boxes in the Information pane.

File

The file menu provides the following options:

- **Open Fabric**—Opens a new switch fabric.
- **Locate Switches and Devices**— Uses the SNMPv2 protocol to discover devices responding to SNMP requests with the read-only community string public. You may use this feature if you want to locate other Cisco MDS 9000 switches in the subnet, but are not physically connected to the fabric.
- **Rediscover**—Initiates an on-demand discovery to learn recent changes from the switches and update the Cisco DCNM-SAN Client. You may use this option when Cisco DCNM-SAN Server is not in sync with switches in the fabric and you do not want to wait until the next polling cycle. The rediscover option does not delete the fabric and add it again. You may delete and add the fabric only if the rediscover option fails to update Cisco DCNM-SAN Server.
- **Resync All Open Fabrics**— Cisco DCNM-SAN Server forces all the fabrics to close and re-open. You may use this option when Cisco DCNM-SAN Client is not in sync with Cisco DCNM-SAN Server.
- **Rediscover SCSI Targets**— Initiates an on-demand discovery to learn recent changes from the SCSI target switches. You may use this option when Cisco DCNM-SAN Server is not in sync with SCSI target switches in the fabric and you do not want to wait until the next polling cycle.
- **Preferences**—Sets your preferences to customize the behavior of the Cisco DCNM-SAN Client.
- **Import Enclosures**—Imports saved enclosures.
- **Export**
 - **Map Image**—Generates and export the map to a specified location.
 - **Visio**—Exports the map to a Visio file.

- Table—Exports the table data to a text file.
- Log—Exports the log to a text file.
- Events—Exports the events to a text file.
- Enclosures—Exports the enclosure values to a text file.
- Print —Prints the map.
- Exit—Exit Cisco DCNM-SAN.

View

View menu provides the following options:

- Refresh Map—Refreshes the current map.
- Layout
 - Cancel—Cancels the current layout.
 - Spring—Displays the layout based on spring algorithm.
 - Quick—Quickly displays the layout when the switch has many end devices.
- Zoom
 - In—Zooms in the view.
 - Out—Zooms out the view.
 - Fit—Fits the view in the fabric pane.
- Grid—Enables the grid view.
- Overview Window—Allows you to center the Fabric pane on the area of the fabric that you want to see. This option is useful for large fabrics that cannot be displayed entirely within the Fabric pane.
- Legend—Shows all the legends used in the fabric map.
- Find in Map—Finds a device in the fabric map.

Zone

The zone menu provides the following options:

- Edit Local Full Zone Database—Allows you to create zones across multiple switches. Zones provide a mechanism for specifying access control. Zone sets are a group of zones to enforce access control in the fabric. All zoning features are available through the Edit Local Full Zone Database dialog box.
- Deactivate Zoneset—Deactivates an active zone set.
- Copy Full Zone Database—Creates a new zone set. On the Cisco MDS Family switches, you cannot edit an active zone set. However, you can copy an active zone set to create a new zone set that you can edit.
- Merge Analysis—Enables you to determine if zones will merge successfully when two Cisco MDS switches are interconnected. If the interconnected switch ports allow VSANs with identical names or contain zones with identical names, then Cisco DCNM-SAN verifies that the zones contain identical members. You can use merge analysis tool before attempting a merge, or after fabrics are interconnected to determine zone merge failure causes.
- Merge Fail Recovery—Recovers the port from its isolated state either by importing the neighboring switch's active zone set database and replacing the current active or by exporting the current database to the neighboring switch.
- Migrate Non-MDS Database—Migrate a non-MDS database using Cisco DCNM-SAN (you may need to use the Zone Migration Wizard to accomplish this task).
- IVR
 - Deactivate Zoneset—Deactivates an active zone set.

- Copy Full Zone Database—Recovers an IVR zone database by copying the IVR full zone database from another switch.
- Copy Full Topology—Recovers a topology by copying from the active zone database or the full zone database.

Tools

Tools menu provides the following options:

- Health
 - Switch Health—Determines the status of the components of a specific switch.
 - Fabric Configuration—Analyzes 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.
 - Show Tech Support—Collects large amount of information about your switch for troubleshooting purposes. When you issue a show tech support command from Cisco DCNM-SAN 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 Cisco DCNM-SAN.
- Connectivity
 - End to End Connectivity—Determines connectivity and routes among devices with the switch fabric. This 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.
 - Ping—Determines connectivity from another switch to a port on your switch.
 - Trace Route—Verifies connectivity between two end devices that are currently selected on the Fabric pane.
 - Compact Flash Report—Automatically scans the fabric and generate a report that shows the status of CompactFlash.
- NPV
 - CFS Static Peer Setup—Manage the peer list used during CFS on NPV-enabled switches. After setting up the static peers list, the CFS discovery on the switches will be changed to static mode for all peers in the list. Cisco DCNM-SAN does not automatically update static peers list. You may need to update the list using the CFS Static Peer Setup Wizard when a new switch is added to the fabric.
 - Traffic Map Setup—Configures the list of external interfaces to the servers, and enabling or disabling disruptive load balancing. Using Traffic Map Setup you can specify the external ports that a server should use for traffic management.
 - Flex Attach Pre-Configure Server—Sets the port configurations for all the ports in a switch such as enabling or disabling FlexAttach, setting the default VSAN ID, and setting the interface status.
 - Flex Attach Move Server—Moves a server to another port on the same NPV device or another NPV device without changing the SAN.
 - Flex Attach Replace Server—Replaces a failed server with a new server on the same port without changing the SAN.
- Data Mobility Manager
 - Server Based—Performs server-based data migration.
 - Storage based—Performs storage-based data migration.

- Server LUN Discovery—Performs LUN discovery to select the LUNs available for migration and automates the session creation by matching the LUNs in the existing and new storage.
- FCoE—Launches the FCoE Configuration Wizard to create virtual Fibre Channel interfaces.
- Port Channel—Creates PortChannels from selected ISL either manually or automatically.
- DPVM Setup—Establishes dynamic port VSAN membership, enables autolearning, and activates the DPVM database.
- IP SAN
 - FCIP Tunnel—Creates FCIP links between Gigabit Ethernet ports. Enables Fibre Channel write acceleration and IP compression.
 - iSCSI Setup—Creates zones for iSCSI initiators and adds a VSAN to a target-allowed VSAN list.
 - SAN Extension Tuner—Optimizes FCIP performance by generating either direct access (magnetic disk) or sequential access (magnetic tape) SCSI I/O commands and directing such traffic to a specific virtual target. This option is used to generate SCSI I/O commands (read and write) to the virtual target based on your configured options.
- Security
 - Port Security—Prevents unauthorized access to a switch port in the Cisco MDS 9000 Family, rejects intrusion attempts and reports these intrusions to the administrator.
 - IP ACL—Creates an ordered list of IP filters in a named IPv4-ACL or IPv6-ACL profile using the IPv4-ACL Wizard.
- Install
 - License—Facilitate download and installation of licenses in selected switches in the fabric.
 - Software—Verifies image compatibility and installs software images on selected switches in the fabric.
- Flow Load Balance Calculator —Allows you to get the best load-balancing configuration for your FICON flows. The calculator does not rely on any switch or flow discovery in the fabric.
- Device Manager—Invokes Device Manager for a switch.
- Command Line Interface —Enables command-line operations.
- Run CLI Commands—Runs command-line operations on more than one switch at a time.

Performance

The performance menu provides the following options:

- Create Flows—Creates host-to-storage, storage-to-host, or bidirectional flows. You can add these flows to a collection configuration file to monitor the traffic between a host or storage element pair.

Server

The server menu provides the following options:

- Admin—Opens the control panel.
- Purge Down Elements—Purges all down elements in the fabric.

Help

The help menu provides the following options:







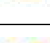

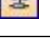





- Contents —Launches the online help contents.
- Config Guide—Launches the Cisco DCNM-SAN Configuration Guide.









- About—Displays information about Cisco DCNM-SAN.

Toolbar

The Cisco DCNM-SAN main toolbar provides icons for accessing the most commonly used menu bar options as shown in [Table 3: Cisco DCNM-SAN Client Main Toolbar](#), on page 201.

Table 3: Cisco DCNM-SAN Client Main Toolbar

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

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

Logical Domains Pane

Use the Logical Domains pane to manage attributes for fabrics, VSANs, and zones, and to access user-defined groups. Starting from NX-OS Release 4.2(0), SAN and LAN nodes are listed under Datacenter node and all the fabrics are listed under SAN node. When you select Datacenter node in the tree, Cisco DCNM-SAN displays all the switches and ISLs. When you select LAN node, Cisco DCNM-SAN displays only Ethernet switches and Ethernet links. Under the fabric node, VSANs are ordered by a VSAN ID. The segmented VSANs are placed under the fabric node. The label next to the segmented VSAN indicates the number of segments. You can expand a segmented VSAN and the segments under that VSAN. Right-click one of the folders in the tree and click a menu item from the pop-up menu. You see the appropriate configuration dialog box.

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

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

You can change the fabric name using Cisco DCNM-SAN.

Procedure

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

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

To further narrow the scope, select attributes from the Physical Attributes pane. The Cisco DCNM-SAN table, display, and filter criteria change accordingly.

Physical Attributes Pane

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

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

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



Note You cannot view the detailed physical attributes of the data center switches or monitor the connections. When you select either a data center node or a LAN node the physical attributes pane will be blank.

Context Menu for Tables

When you right-click in the table, you see a pop-up menu with options that vary depending on the type of option you selected in the Physical Attributes pane. You can perform various operations by right-clicking the device listed in the table. To view various options available for switches, ISLs, and end devices, refer to the procedures in the sections that follows:

Viewing Switch Options

When you select the datacenter node, the switch table displays all the switches that are discovered. When you select the SAN node or the fabric node, the switch table displays all the Fibre Channel switches and when you select the LAN node, the switch table displays all the Ethernet switches.

Procedure

Step 1 Click Switches in the Physical Attributes pane.

Step 2 Right-click on the device in the table.

The pop-up menu provides the following options:

- Apply Changes—Applies the changes to the switch.
- Refresh Values—Refreshes the current values.
- Undo Changes—Undoes modifications to the switch.
- Export to File—Export the values to a file.
- Print Table—Prints the table.
- Detach Table—Detaches the table.
- Switch Attributes—Changes the switch properties.
- Interface Attributes—Changes the interface properties.
- Element Manager—Manages this switch.
- Command Line Interface—Enables to perform command line operations.
- Copy—Copies the switch.
- Purge—Purges the switch.
- Fix Location—Fixes the switch in the current location.
- Align—Aligns the switch.
- Show End Devices—Shows the end devices.
- Expand Multiple Links—Expands the links to this switch.
- Other—Other options
- Group—Groups switches

Viewing ISL Options

When you select the data center node, the ISLs table displays all of the Fibre Channel and Ethernet links. When you select the LAN node, the ISLs table displays all the Ethernet links.

Procedure

Step 1 In the Physical Attributes pane, click ISLs and then click Summary tab.

Step 2 Right-click the device in the table.

The pop-up menu provides the following options:

- Refresh Values—Refreshes the current values.
- Copy—Copies information from a specific field.
- Find—Conducts search based on the input string.
- Export to File—Exports the values to a file.
- Print Table—Prints the table.
- Detach Table—Detaches the table.
- Interface Attributes—Changes the interface properties.
- Element Manager—Manages the device.
- FCIP Tunnel Attributes—Changes FCIP tunneling properties.
- Create Port Channel—Creates port channel.
- Re-enable—Reenables a disabled device.
- Enable FC-SP—Enables FC-SP.
- SAN Extention Tuner—Optimizes FCIP performance.
- Purge—Purges the device.

When you select a port channel from the table, the pop-up menu will have the following additional options:

- Member Attributes—Changes the member properties.
- Channel Attributes—Changes the port channel properties.
- Edit—Edits the channel properties.

Viewing End Device Options

Procedure

Step 1 In the Physical Attributes pane, click End Devices and then click the Summary tab.

Step 2 Right-click the device in the table.

The pop-up menu provides the following options:











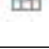
- Apply Changes—Applies the changes to the device.
- Refresh Values—Refreshes the current values.
- Copy—Copies the information specific to the field.
- Paste—Pastes the copied text.
- Undo Changes—Undoes modifications to the device.
- Find—Searches for information depending on the input string.
- Export to File—Exports the values to a file.
- Print Table—Prints the table.
- Detach Table—Detaches the table.
- Device Attributes—Changes the device properties.
- Interface Attributes—Changes the interface properties.
- Element Manager—Manages this device.
- Command Line Interface—Enables you to perform command line operations.
- Copy—Copies the switch.
- Purge—Purges the switch.
- Fix Location—Fixes the switch in the current location.

- Align—Aligns the switch.
- Ping—Pings another device.
- Trace Route—Determines the route taken by packets across the network.
- Select Dependent Ports—Selects dependent ports.
- Group—Groups devices.

Information Pane

Use the Information pane to display tables of information associated with the option selected from the menu tree in the Logical Domains or Physical Attributes panes. The Information pane toolbar provides buttons for performing one or more of the operations shown in [Table 4: Information Pane Toolbar](#), on page 206.

Table 4: Information Pane Toolbar

Icon	Description
	Applies configuration changes.
	Refreshes table values.
	Opens the appropriate dialog box to make a new row in the table.
	Deletes the currently highlighted rows from the table.
	Copies data from one row to another.
	Pastes the data from one row to another.
	Undoes the most recent change.
	Finds a specified string in the table.
	Exports and saves information to a file.
	Prints the contents of the Information pane.
	Displays a non-editable copy of the table in the Information pane in its own window, which you can move around the screen.



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



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









Detachable Tables






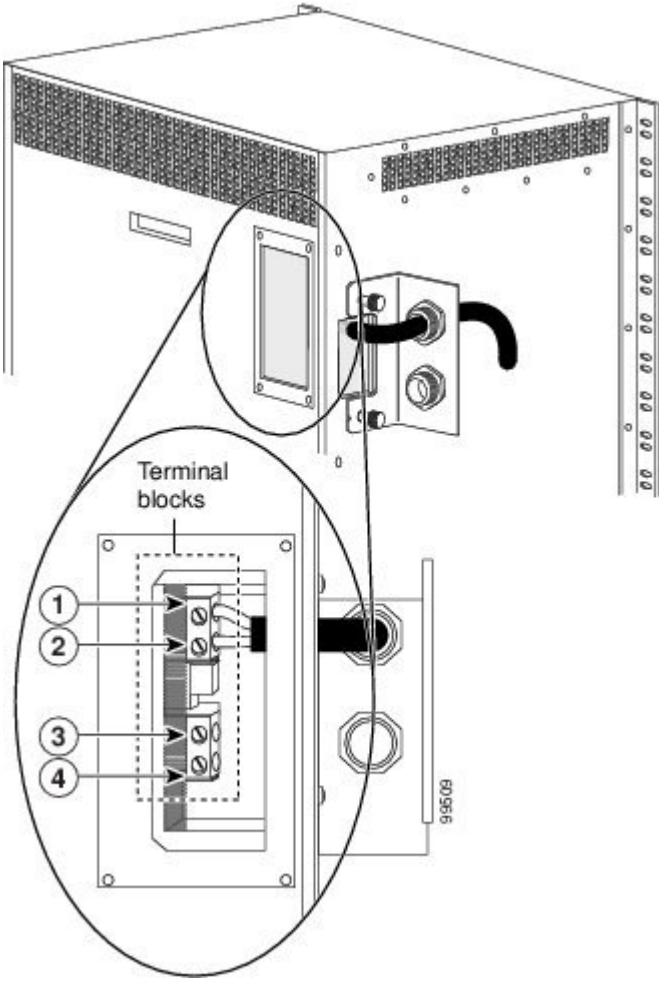

Detachable tables in Cisco DCNM-SAN allow you to detach tables and move them to different areas on your desktop so that you can compare similar tables from different VSANs. You can keep informational tables open from one view while you examine a different area in Cisco DCNM-SAN. To detach tables, click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




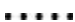




Fabric Pane

Use the Fabric pane to display the graphical representation of your fabric. [Table 5: Cisco DCNM-SAN Graphics](#), on page 207 explains the graphics you may see displayed, depending on which devices you have in your fabric.

Table 5: Cisco DCNM-SAN Graphics

Icon or Graphic	Description
	Director class MDS 9000 Fibre Channel switch.
	Non-director class MDS 9000 Fibre Channel switch.
	Nexus 7000 switch.
	Nexus FCoE or Fibre Channel switch.
	Catalyst LAN switch.
	Generic Fibre Channel switch.
	Cisco SN5428.
	Dashed or dotted orange line through a device indicates that the device is manageable but there are operational problems.

Icon or Graphic	Description
	Dashed or dotted orange X through a device or link indicates that the device or ISL is not working properly.
	A red line through a device indicates that the device is not manageable.
	A red X through a device or link indicates that the device is down or that the ISL is down.
	Fibre Channel HBA (or enclosure).
	Fibre Channel target (or enclosure).
	iSCSI host.
	Fibre Channel ISL and edge connection.

Icon or Graphic	Description
	Fibre Channel PortChannel.
	IP ISL and edge connection.
	IP PortChannel.
	DWDM connection.
	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 Cisco DCNM-SAN 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, Cisco DCNM-SAN can no longer communicate with it.

The bottom of the Fabric pane has the following tabs:

- Fabric—When displaying multiple fabrics, each fabric has its own tab. You can switch between fabrics by clicking on their respective tabs.
- Log—Displays messages that describe Cisco DCNM-SAN 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, SNMP, and FICON.

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

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



Note

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

Context Menus

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

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

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

**Note**

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

Saving the Map

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

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

To save the map as a Visio diagram, 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:

- Choose 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 purge it, that element will reappear on the next fabric discovery cycle.

Multiple Fabric Display

Cisco DCNM-SAN can display multiple fabrics in the same pane.

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



Note Enclosure names should be unique. If the same enclosure name is used for each port, Cisco DCNM-SAN 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.

Procedure

-
- Step 1** Right-click a switch or end port in the Fabric pane map and select **Group > Create**.
You see the Edit User Defined Group dialog box.
- Step 2** Enter a group name in the **Name** field.
- Step 3** Use the arrows to move additional switches or end ports from the **Available** column to the **Selected** column.
- Step 4** Click **OK** to save the group.
- Step 5** To add a switch or end port to an existing group in Cisco DCNM-SAN.
- Right-click a switch or end device and select **Group > Add To > YourGroupName**.
You see the Edit User Defined Group dialog box.
 - Use the arrows to move additional switches or end ports from the **Available** column to the **Selected** column.
 - Click **OK** to save the updated group.
- Step 6** To filter the display by a group you have created.
- Expand the **Groups** folder in the Logical Domains pane.
You see the list of groups that you have created.
 - 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.
 - Click the **Groups** folder in the Logical Domains pane to return the display to normal.
- Note** User-defined groups tables are filtered based on switches in the group except for switches where CFS-controlled features are enabled when all CFS member switches are displayed to avoid misconfigurations.
-

Status Bar

The status bar at the bottom of the Cisco DCNM-SAN 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.

Launching Cisco DCNM-SAN Client

As of Cisco SAN-OS 3.x and NX-OS Release 4.x, the Fabric Manager Client login procedure has changed.

Launching Fabric Manager Client in Cisco SAN-OS Release 3.2(1) and Later

You can launch Fabric Manager Client.

**Note**

Network administrators must initially launch Cisco DCNM-SAN Client using Cisco DCNM-SAN Web Server, as described in the following procedure. Once an administrator has installed the Cisco DCNM-SAN Client icon on your desktop, you can double-click the icon to launch the Cisco DCNM-SAN Client.

Procedure

- Step 1** Open your browser and enter the IP address where you installed Cisco DCNM-SAN Server, or enter localhost if you installed Cisco DCNM-SAN Server on your local workstation.
- You see the Cisco DCNM Web Client Login dialog box.
- Step 2** Enter your user name and password and click **Login**.
- You see the Cisco DCNM Web Client Summary page.
- Step 3** Click the **Download** link in the upper right corner of the page.
- You see the Download page for Cisco DCNM-SAN and Device Manager.
- Step 4** Click the link for **Cisco DCNM-SAN**.
- If you are launching Cisco DCNM-SAN Client for the first time, you see a message asking whether you want to create shortcuts for Cisco DCNM-SAN.
- Step 5** Click **Yes** to create shortcuts for Cisco DCNM-SAN.
- Note** This message only appears the first time you launch Cisco DCNM-SAN Client. If you select No, your selection will be remembered and you will not be prompted to make a selection again. In this case, you will need to launch Cisco DCNM-SAN Client using the Cisco DCNM-SAN Web Client.
- Step 6** When the software is installed and icons are created on your desktop, double-click the Cisco DCNM-SAN icon to launch Cisco DCNM-SAN.
- You see the Cisco DCNM-SAN Login dialog box.
- Step 7** Enter the Cisco DCNM-SAN Server user name and password.

- Step 8** Check the Use SNMP Proxy check box if you want Cisco DCNM-SAN Client to communicate with Cisco DCNM-SAN Server through a TCP-based proxy server.
- Step 9** Click **Login**. Once you successfully log in to Cisco DCNM-SAN Server, you can set the seed switch and open the fabrics that you are entitled to access.
- Note** When you launch Cisco DCNM-SAN 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.
- Note** Only network administrators can discover new fabrics.
- Step 10** Click the Ethernet (CDP) radio button to discover using Cisco Discovery Protocol (CDP).
- Step 11** Starting from NX-OS Release 4.2(0), Fabric Manager uses Cisco Discovery Protocol to discover Ethernet switches such as Nexus 5000, Nexus 7000, Catalyst 4000, and Catalyst 6000 switches. You need to use a CDP seed switch for a CDP discovery. Set the fabric seed switch to the Cisco MDS 9000 Family switch or Cisco Nexus 5000 Series that you want Fabric Manager to use.
- Step 12** Choose the Auth-Privacy option according to the privacy protocol you have configured on your switch:
- 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 choice.
- Note** You may use SNMP v2 credentials for CDP discovery as the most of the Catalyst switches do not use MD5-DES for configuration.
- Note** If you want a clean fabric discovery, remove the fabric and rediscover it. If you want a clean LAN discovery, unmanage LAN, remove the CDP seed switch and then rediscover it.
- Step 13** Enter the username and password for the switch.
- Step 14** (Optional) To limit the discovery, specify the VSAN range. Scoping limits the resources discovered by Cisco DCNM-SAN client. You can either include a range of VSANs to be discovered or exclude a range of VSANs from being discovered.
- To limit the discovery to a range of VSANs, click Included VSAN List radio button. Specify the range of VSANs.
 - To exclude a range of VSANs from being discovered, click Excluded VSAN List radio button. Specify the range of VSANs to be excluded.
- Step 15** Click **Discover**.
- You see the Control Panel dialog box.
- You see the included and excluded VSANs list under the Fabric tab.
- 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).
- Note** In the open tab, you see all the discovered fabrics displayed in the control panel. You need to click on the Open button to see all the discovered Ethernet switches.
- Step 16** Check the check box(es) next to the fabric(s) you want to open in the Select column, 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 2-6](#).

Step 17 Click **Open** to open the selected fabric(s).

- Note**
- If you have an incomplete view of your fabric, rediscover the fabric with a user that has no VSAN restriction.
 - If the fabric includes a Cisco Nexus 5000 Series switch, then the Layer 2 node appears under the Switches > Interfaces > Ethernet tree, the VFC (FCoE) node appears under the Switches > Interfaces tree, and the FCoE node appears under the Switches tree in the Physical Attributes pane.
 - For Cisco Nexus 5000 Series switches in the fabric, the tooltip for the switch shows the bind information of a virtual Fibre Channel interface to its corresponding Ethernet interface, such as vfc2(eth1/4).

Launching Fabric Manager Client in Cisco SAN-OS Release 3.2(1) and Later

You can launch Cisco DCNM-SAN Client from within a running instance of Cisco DCNM-SAN.

Procedure

Step 1 Choose **File > Open** or click the **Open Switch Fabric** icon on the Cisco DCNM-SAN toolbar.

You see the Control Panel dialog box.

Step 2 Check the check box(es) next to the fabric(s) you want to open in the Select column and click **Open**.

Note Changes made using Cisco DCNM-SAN 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), Cisco DCNM-SAN prompts you to save your changes before you exit.

Launching Cisco DCNM-SAN Client Using Launch Pad

Starting from Cisco NX-OS Release 4.2(0), you can use Cisco DCNM-SAN launch pad to connect to any server by specifying the IP address of the server. With launch pad, you can connect to any Cisco DCNM-SAN Server version 3.3(0) and later. Launch pad establishes connection with the server using HTTP protocol.

Procedure

Step 1 Open your browser and enter the IP address where you installed Cisco DCNM-SAN Server, or enter localhost if you installed Cisco DCNM-SAN Server on your local workstation.

You see the Cisco DCNM-SAN Web Server Login dialog box.

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

You see the Cisco DCNM-SAN Web Client Summary page.

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

You see the Download page for Cisco DCNM-SAN and Device Manager.

Step 4 Click the link for **Cisco DCNM-SAN**.

You see the Cisco DCNM-SAN Server launch pad.

Step 5 Enter the host name of the server or IP address in the Server URL drop-down list.

Step 6 Click Start.

Note Launch pad retains the history of the server URLs used. You can choose one of the previously user Server URLs from the drop-down list.

Setting Cisco DCNM-SAN Preferences

To set your preferences for the behavior of the Cisco DCNM-SAN, choose **File > Preferences** from the Cisco DCNM-SAN 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 Cisco DCNM-SAN 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 Cisco DCNM-SAN. Check the Prepend Name check box to display the name in front of the IP address of the switch. Check the Replacing Vendor Bytes check box to display the name instead of the IP address. The default is the Prepend Name option.
- **Show End Device Using**—Displays end devices in the Fabric pane using alias or pWWN alias. The default setting for this value is Alias.
- **Show Shortened iSCSI Names**—Displays the default setting for this value is OFF.
- **Show Timestamps as Date/Time**—Displays timestamps in the date/time format. If this preference is not checked, timestamps are displayed as elapsed time. The default setting is enabled (checked).
- **Telnet Path**—Displays the path for the telnet.exe file on your system. The default is **telnet.exe**, but you need to browse for the correct location.

**Note**

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

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

The default SNMP preferences for Cisco DCNM-SAN 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 ON.
- **Enable Audible Alert when Event Received**—The default setting for this value is OFF.

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

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

Network LAN Discovery

Starting from NX-OS Release 4.2(0), you can discover Nexus and Catalyst Ethernet switches using Cisco Discovery Protocol (CDP). DataCenter 3(DC3) switches are displayed under Datacenter and LAN nodes. Cisco DCNM-SAN displays basic information about DC3 switches and its ISLs.

Viewing Ethernet Switches

Procedure

Step 1 Click the LAN node under Datacenter node.

Step 2 Click Switches tab in the Information pane.

You can see the switch information.

Note Datacenter is the parent node of SAN and LAN nodes. The SAN node remains in the tree as the parent for all the fabrics.

Removing a LAN

Procedure

- Step 1** Choose Server > Admin.
You can see the switch information.
- Step 2** Click to select the switch IP of the LAN you want to remove.
- Step 3** Click Remove.
-

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 Cisco DCNM-SAN map.

Procedure

- Step 1** Expand **End Devices** and then choose **Storage** or **Hosts** in the Physical Attributes pane.
You see the end devices displayed in the Information pane.
- Step 2** Click one of the devices in the Fabric pane, or click the **Enclosures** tab of the Information pane, 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 select more than one name, press the Shift key and click each of the other names.
- Step 5** Press **Ctrl-C** to copy the selected name(s).
- Step 6** Press **Ctrl-V** to paste the device name into the Name field.

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

Using Alias Names as Enclosures

Procedure

- Step 1** Expand End Devices and choose Hosts or Storage from the Physical Attributes pane.
You see the list of devices in the Information pane. The NxPorts tab is the default.
- Step 2** Right-click the enclosure names that you want to convert to alias names and choose Alias > Enclosure .

The Alias > Enclosures window appears . It contains a list of expressions. You can also add expressions to the list and modify expressions in the current list.

Step 3 Click the Apply Changes icon to save the changes and then click Close.

Note Cisco DCNM-SAN uses the regular expressions to convert multiple alias names into one enclosure. The alias names should be in the same expression pattern rule. You can create enclosure names from selected aliases using the regular expressions list.

Using Alias Names as Descriptions

Procedure

Step 1 Choose End Devices and from the Physical Attributes pane.

Step 2 Click the General tab.

You see the list of devices in the Information pane.

Step 3 Select the device names that you want to populate the description with alias names and then click Alias > Enclosure button.

You see the alias names are copied to corresponding rows in the description column.

Note Cisco DCNM-SAN does not parse or format the alias name while copying.

Controlling Administrator Access with Users and Roles

Cisco MDS 9000 Family switches support role-based management access whether using the CLI or Cisco Cisco DCNM-SAN. This 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 (Cisco DCNM-SAN and Device Manager). These access permissions are automatically granted to all users in order for them to use the GUI.

Cisco Cisco DCNM-SAN 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 Cisco DCNM-SAN to create roles and users and to assign passwords as required for secure management access in your network.



Note Either to create a new SNMPv3 user or modify password of SNMPv3 user, the DCNM login user need to have enabled with DES/AES privacy password. Since the creating and modifying SNMP SET request need to be encrypted, the login user password needs to have the privacy password.

Using Cisco DCNM-SAN Wizards

Cisco DCNM-SAN Client provides the following wizards to facilitate common configuration tasks:

- **VSAN**—Creates VSANs on multiple switches in the fabric and sets VSAN attributes including interop mode, load balancing, and FICON.
- **Zone Edit Tool**—Creates zone sets, zones, and aliases. Adds members to zones and edits the zone database.
- **IVR Zone**—Creates IVR zone sets, zones, and aliases. Enables IVR NAT and auto-topology. Adds members to IVR zones, and edits the IVR zone database.
- **FCoE**—Creates virtual Fibre Channel (FC) interfaces and VLAN-VSAN mappings, and binds virtual FC interfaces to Ethernet interfaces or PortChannels.
- **PortChannel**—Creates PortChannels from selected ISLs either manually or automatically. Sets PortChannel attributes such as channel ID and trunking mode.
- **FCIP**—Creates FCIP links between Gigabit Ethernet ports. Enables Fibre Channel write acceleration and IP compression.
- **DPVM**—Establishes dynamic port VSAN membership, enables autolearning, and activates the DPVM database.
- **Port Security**—Prevents unauthorized access to Cisco MDS switches and reports these intrusions to the administrator.
- **iSCSI**—Creates zones for iSCSI initiators and adds a VSAN to a target-allowed VSAN list.
- **NPV**—Reduces the number of Fibre Channel domain IDs in SANs.
- **QoS**—Sets QoS attributes for zones in the selected VSAN.
- **IP ACL**—Creates ordered IP access control lists and distributes to selected switches in the fabric.
- **License Install**—Facilitates download and installation of licenses in selected switches in the fabric.
- **Software Install**—Verifies image compatibility and installs software images on selected switches in the fabric.

Cisco DCNM-SAN Troubleshooting Tools

Cisco DCNM-SAN has several troubleshooting tools available from the toolbar or Tools menu

- **Zone Merge Analysis**

—The zone merge analysis tool (available from the Zone menu) enables you to determine if zones will merge successfully when two Cisco MDS switches are interconnected. If the interconnected switch ports allow VSANs with identical names or contain zones with identical names, then Cisco DCNM-SAN 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**—Cisco DCNM-SAN's end-to-end connectivity analysis tool uses FC Ping to verify interconnections between Cisco MDS switches and end-device (HBAs and storage devices) in a particular VSAN. In addition to basic connectivity, Cisco DCNM-SAN 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 Cisco DCNM-SAN. It verifies the status of all critical Cisco MDS switches, modules, ports, and Fibre Channel services. Over 40 conditions are checked. This tool provides a very fast, simple, and thorough way to assess Cisco MDS switch health.
- **Fabric Configuration Analysis**—Cisco DCNM-SAN includes a fabric configuration analysis tool. It compares the configurations of all Cisco MDS switches in a fabric to a reference switch or a policy file. You can define what functions to check and what type of checks to perform. The analysis can look for mismatched values, and missing or extra values. If all configuration checking is performed for all functions, over 200 checks are performed for each Cisco MDS switch.

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

Integrating Cisco DCNM-SAN and Data Center Network Management Software

Cisco DCNM-SAN and Data Center Network Management (DCNM) software are the two major components in the Cisco next-generation data center environment. Cisco DCNM-SAN configures Cisco Nexus 5000 Series switches and Cisco MDS 9000 Series switches. DCNM software configures Cisco Nexus 5000 and Cisco Nexus 7000 Series switches. The Scope of the Cisco DCNM-SAN software is confined to SAN while the scope of the DCNM-LAN software is limited to the LAN network.

In a typical data center environment, the mixture of SAN and LAN topology are becoming increasingly common. Since the two management software are not designed to work across their topology limits, users are not able to navigate to Cisco DCNM-SAN from DCNM-LAN software and vice versa.

Integrating Cisco DCNM-SAN and DCNM-LAN provides a single platform to manage the networks in data center 3.0 and it provides seamless user experience under specific configuration. Starting from Cisco MDS NX-OS Release 4.2, the directory structure has changed to accommodate the integration of Cisco DCNM-SAN with Cisco Nexus 5000 Series products.

Launching a Switch from the Topology Map

Procedure

-
- | | |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | In the Cisco DCNM-SAN fabric pane, right-click the Nexus switch in the LAN map that you want to open with DCNM.

You see the pop-up menu. |
| Step 2 | In the pop up menu, click DCNM and select appropriate context. |
-



CHAPTER 5

Configuring Device Manager

- [Device Manager](#), on page 223

Device Manager

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

Information About Device Manager

Device Manager provides a graphic representation of a Cisco MDS 9000 Family switch chassis or Cisco Nexus 5000 Series switch chassis, or a Cisco Nexus 7000 Series switch chassis including the installed switching modules, the supervisor modules, the status of each port within each module, the power supplies, and the fan assemblies.



Note

Device Manager support for Cisco Nexus 7000 Series switches is only for FCoE. Non-FCoE modules appear as Unsupported Card.

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

Device Manager Release 4.2 and later provides enhanced security using multiple perspectives (simple and advanced) allowing role based-access to its features. The Device Manager perspective filters out menu items that are not relevant to the user. Users with server admin role, can only access a subset of the fabric related features. The server admin role will not be able to manage Device Manager users or connected clients.

Device Manager Release 5.0 and later supports all the software features that are offered by Cisco NX-OS for managing Cisco MDS 9148 and 9124 Multilayer Fabric switches. Cisco MDS 9148 Multilayer Fabric Switch is a 48-port (1/2/4/8G) FC 1RU switch based on the Sabre ASIC and Cisco MDS 9124 Multilayer Fabric switch is a 1/2/4/8G switch module for HP BladeServer based on the Sabre ASIC. Device Manager and DCNM-SAN allow you to discover, display, configure, monitor and service both these new switches. Device Manager also supports the following Cisco Nexus 2000 Series Fabric Extenders on a Cisco Nexus 5000 Series switch that runs Cisco NX-OS Release 5.0(1):

- Cisco Nexus 2148T Fabric Extender—It has four 10-Gigabit Ethernet fabric interfaces for its uplink connection to the parent Cisco Nexus 5000 Series switch and eight 1-Gigabit Ethernet or 10-Gigabit Ethernet host interfaces for its downlink connection to servers or hosts.
- Cisco Nexus 2232PP Fabric Extender—It has eight 10-Gigabit Ethernet fabric interfaces with SFP+ interface adapters for its uplink connection to the parent Cisco Nexus 5000 Series switch and 32 10-Gigabit Ethernet fabric interfaces with SFP+ interface adapters for its downlink connection to servers or hosts.
- Cisco Nexus 2248TP Fabric Extender—It has four 10-Gigabit Ethernet fabric interfaces with small form-factor pluggable (SFP+) interface adapters for its uplink connection to the parent Cisco Nexus 5000 Series switch and 48 1000BASE-T (1-Gigabit) Ethernet host interfaces for its downlink connection to servers or hosts.

Device Manager allows you to discover and display these Fabric Extenders. Cisco Device Manager and the Cisco DCNM-SAN client support provisioning and monitoring of the 48-port 8-Gbps Advanced Fibre Channel switching module (DS-X9248-256K9) and the 32-port 8-Gbps Advanced Fibre Channel switching module (DS-X9232-256K9).

Device Manager Features

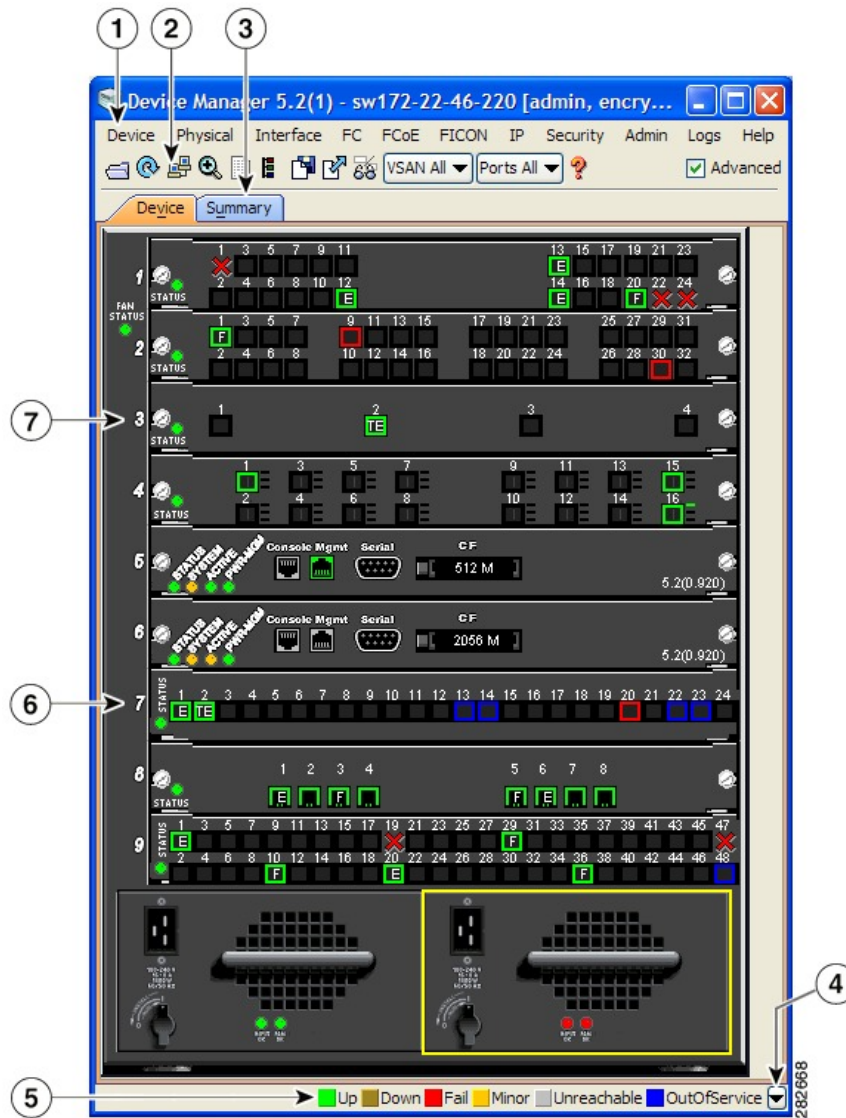
Device Manager provides two views: Device View and Summary View. Use Summary View to monitor interfaces on the switch. Use Device View to perform switch-level configurations including the following:

- Configure virtual Fibre Channel interfaces.
- Configure Fibre Channel over Ethernet (FCoE).
- Configure zones for multiple VSANs.
- Manage ports, PortChannels, and trunking.
- Manage SNMPv3 security access to switches.
- Manage CLI security access to the switch.
- Manage alarms, events, and notifications.
- Save and copy configuration files and software image.
- View hardware configuration.
- View chassis, module, port status, and statistics.

Using Device Manager Interface

This section describes the Device Manager interface.

Figure 3: Device Manager, Device Tab



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

Menu Bar

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

- **Device**—Opens an instance of Device Manager, sets management preferences, sets the page layout, opens a Telnet/SSH session with the current switch, exports a device image, and closes the Device Manager application.
- **Physical**—Allows you to view and manage inventory, modules, temperature sensors, power supplies, fans, and the entire system.
- **Interface**—Allows you to configure and manage PortChannels, as well as Fibre Channel, Ethernet, iSCSI, and FICON ports. Also provides diagnostic, management and monitoring capabilities, as well as SPAN and port tracking.

**Note**

The Interface > Port Channels menu option does not appear if the Cisco Nexus 5000 Series switch is in NPV mode and runs a Cisco NX-OS release prior to 4.2(1).

- **FC**—Allows you to configure and manage VSAN, domain, and name server characteristics. Also provides advanced configuration capabilities.
- **FCoE**—Allows you to configure the FCoE parameters and map VSANs to VLANs on a Cisco Nexus 5000 Series switch.

**Note**











The FCoE menu option appears only if the Cisco Nexus 5000 Series switch runs Cisco NX-OS Release 4.0(1a) or later releases.


- **FICON**—Allows you to configure and manage FICON VSANs, configure RLIR ERL information, swap selected FICON ports, and view FICON port numbers.
- **IP**—Allows you to configure and manage the following types of information: FCIP, iSCSI, iSNS, routes, VRRP, and CDP.
- **Security**—Allows you to configure and manage FCSP, port security, iSCSI security, SNMP security, common roles, SSH, AAA, and IP ACLs.
- **Admin**—Allows you to save, copy, edit, and erase the switch configuration, monitor events, manipulate Flash files, manage licenses, configure NTP, use CFS, and reset the switch. Also enables you to use the show tech support, show cores, and show image commands.
- **Logs**—Shows the various logs: message, hardware, events, and accounting. Also displays FICON link incidents, and allows you to configure the syslog setup.
- **Help**—Displays online help topics for specific dialog boxes in the Information pane.

Toolbar Icons

The Device Manager toolbar provides quick access to many Device Manager features. Once the icon is selected, a dialog box may open that allows configuration of the feature. The toolbar provides the main Device and Summary View icons as shown in [Table 6: Device Manager Main Toolbar](#), on page 227.

Table 6: 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.
 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.
 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.

Icon	Description
 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 for descriptions of these icons.

Tabs

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



Note The Device view also shows the switch chassis information of the Cisco Nexus 2000 Series Fabric Extenders (FEXs) that are connected to a Cisco Nexus 5000 Series switch that runs Cisco NX-OS Release 5.0(1).

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



Note The Summary tab does not display the utilization statistics (Util%) of virtual Fibre Channel interfaces for Cisco Nexus 5000 Series switches that run Cisco NX-OS Release 4.2.

Legend

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

Colors

- Green—The port is up.
- Brown—The port is administratively down.
- Red—The port is down or has failed.
- Amber—The port has a minor fault condition.
- Gray—The port is unreachable.
- 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)
- TNP—Trunking NP_Port (NPV Mode)
- TF—Trunking F_Port
- f—vFC Present (Cisco Nexus 5000 Series switches only)

Supervisor and Switching Modules

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



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

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

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

To create PortChannels using Device Manager, click **PortChannels** from the Interface menu.



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

Context Menus

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

From Device View:

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

From Summary View:

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

Launching Device Manager

To launch Device Manager from your desktop, double-click the **Device Manager** icon and follow the instructions described in the Cisco DCNM Installation and Licensing Guide.

Procedure

-
- Step 1** You can choose one of the following three steps
- 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**.
- You see the Device Manager open dialog box.
- Step 2** Enter the IP address of the device.
- Step 3** Enter the user name and password.
- Step 4** Check the Proxy SNMP through FMS check box if you want Device Manager Client to use a TCP-based proxy server.
- Step 5** Choose the Auth-Privacy option according to the privacy protocol you have configured on your switch:
- 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 choice.
- Step 6** Click Open to open the Device Manager.
-

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).
- **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**").

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



CHAPTER 6

Configuring Performance Manager

- [Configuring Performance Manager, on page 233](#)

Configuring Performance Manager

This chapter describes how DCNM-SAN is used to monitor and manage a network. This chapter includes the following topics:

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 DCNM-SAN 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 DCNM-SAN Server.

Data Interpolation

One of the unique features of Performance Manager is its ability to interpolate data when statistical polling results are missing or delayed. Other performance tools may store the missing data point as zero, but this can distort historical trending. Performance Manager interpolates the missing data point by comparing the data point that preceded the missing data and the data point stored in the polling interval after the missing data. This maintains the continuity of the performance information.

Data Collection

One year's worth of data for two variables (Rx and Tx bytes) requires a round-robin database (rrd) file size of 76 K. If errors and discards are also collected, the rrd file size becomes 110 K. The default internal values are as follows:

- 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 worth of historical data that includes errors and discards. If there were 20 switches in this SAN with equal distribution of fabric ports, about two to three SNMP packets per switch would be sent every 5 minutes for a total of about 100 request or response SNMP packets required to monitor the data.

Because of their variable counter requests, flows are more difficult to predict storage space requirements for. But in general you can expect that, each extra flow adds another 76 KB.



Note Performance Manager does not collect statistics on nonmanageable and non-MDS switches. Loop devices (FL/NL) are not collected.

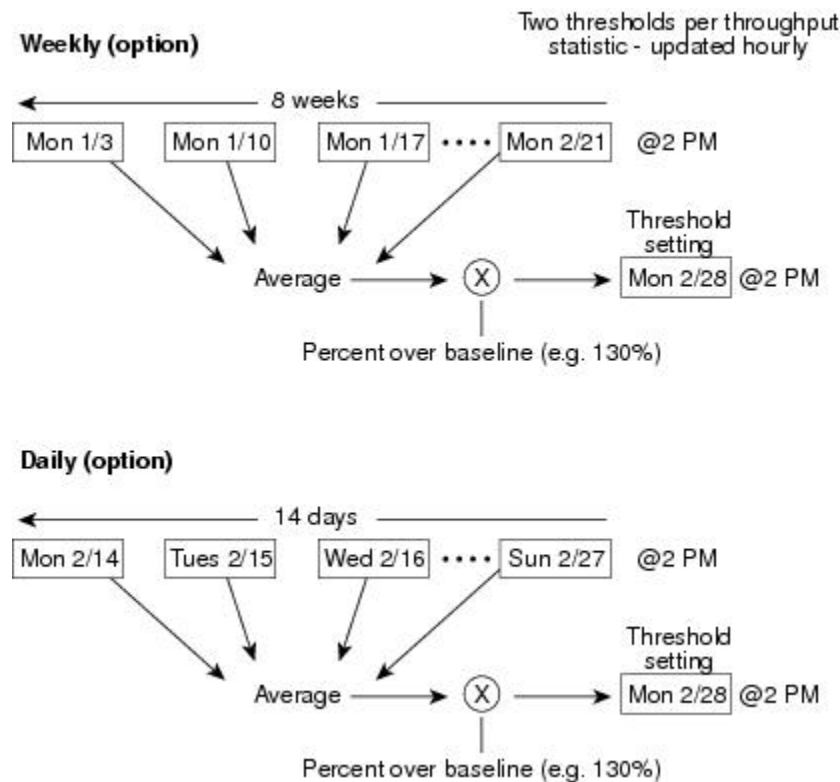
Using Performance Thresholds

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 DCNM-SAN 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 DCNM-SAN web client Events tab.

Baseline thresholds create a threshold that adapts to the typical traffic pattern for each link for the same time window each day, week, or every two weeks. Baseline thresholds are set as a percent of the average (110% to 500%), where 100% equals the calculated weighted average. [Figure 4: Baseline Threshold Example, on page 235](#) shows an example of setting a baseline threshold for a weekly or daily option.

Figure 4: Baseline Threshold Example



The threshold is set for Monday at 2 p.m. The baseline threshold is set at 130% of the average for that statistic. The average is calculated from the statistics value that occurred at 2 p.m. on Monday, for every prior Monday (for the weekly option) or the statistics value that occurred at 2 p.m. on each day, for every prior day (for the daily option).

Flow Statistics

Flow statistics count the ingress traffic in the aggregated statistics table. You can collect two kinds of statistics:

- Aggregated flow statistics to count the traffic for a VSAN.
- Flow statistics to count the traffic for a source and destination ID pair in a VSAN.

If you enable flow counters, you can enable a maximum of 1 K 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.

Generation 1 modules allow a maximum of 1024 flow statements per module. Generation 2 modules allow a maximum of 2048-128 flow statements per module.

[Table 7: Performance Manager Flow Types, on page 236](#) explains the Flow Type radio button that defines the type of traffic monitored.

Table 7: Performance Manager Flow Types

Flow type	Description
Host->Storage	Unidirectional flow, monitoring data from the host to the storage element
Storage->Host	Unidirectional flow, monitoring data from the storage element to the host
Both	Bidirectional flow, monitoring data to and from the host and storage elements

Flow Setup Wizards

The Performance Manager Flow and Performance Manager Setup wizards greatly simplify configuration. All you need to do is select the categories of statistics to capture and the wizards provide a list of flows and links to monitor. You can remove entries if desired, or just accept the provided list and start data collection. Statistics for host and storage links are not associated with a specific port on a switch, so you do not lose long term statistics if a connection is moved to a different port.

Creating a Flow Using Performance Manager Flow Wizard

Procedure

- Step 1** Choose Performance > Create Flows.
- You see the Define Traffic Flows dialog box.
- Step 2** Click the drop-down menu in the VSAN field.
- Step 3** Choose the list of VSANs provided by the flow configuration wizard.
- Step 4** Click the drop-down menu in the Zone field.
- Step 5** Choose the list of zones provided by the flow configuration wizard.
- Step 6** Click Next to continue to the next window.
- Step 7** Choose items in the Possible Flow Pairs area.
- The Review Traffic Flows window displays all VSAN flow pairs in the Existing Flows for Vsan area.
- Step 8** Click Add to create the selected flow.
- Step 9** Choose items in the Existing Flows for Vsan area.
- Step 10** Click Remove to remove the selected flow.
- Step 11** Click Finish to restart the Performance Manager collection.
- You see the Confirmation dialog box.
- To verify the newly created flow, choose Physical Attributes > End Devices > Flow Statistics. The newly created flows are displayed.
-



CHAPTER 7

Configuring High Availability

- [Configuring High Availability, on page 237](#)

Configuring High Availability

This chapter describes how to configure high availability, and describes the switchover processes.

About High Availability

Process restartability provides the high availability functionality in Cisco MDS 9000 Family switches. This process ensures that process-level failures do not cause system-level failures. It also restarts the failed processes automatically. This process is able to restore its state prior to the failure and continues executing from the failure point going forward.

An HA switchover has the following characteristics:

- It is stateful (nondisruptive) because control traffic is not impacted.
- It does not disrupt data traffic because the switching modules are not impacted.
- Switching modules are not reset.



Note Switchover is not allowed if auto-copy is in progress.

Switchover Processes

Switchovers occur by one of the following two processes:

- The active supervisor module fails and the standby supervisor module automatically takes over.
- You manually initiate a switchover from an active supervisor module to a standby supervisor module.

Once a switchover process has started another switchover process cannot be started on the same switch until a stable standby supervisor module is available.

**Caution**

If the standby supervisor module is not in a stable state (ha-standby), a switchover is not performed.

This section includes the following topics:

Synchronizing Supervisor Modules

The running image is automatically synchronized in the standby supervisor module by the active supervisor module. The boot variables are synchronized during this process.

The standby supervisor module automatically synchronizes its image with the running image on the active supervisor module.

**Note**

The image a supervisor module is booted up from cannot be deleted from bootflash. This is to ensure that the new standby supervisor module is able to synchronize during the process.

Manual Switchover Guidelines

Be aware of the following guidelines when performing a manual switchover:

- When you manually initiate a switchover, system messages indicate the presence of two supervisor modules.
- A switchover can only be performed when two supervisor modules are functioning in the switch.
- The modules in the chassis are functioning as designed.

Manually Initiating a Switchover

To perform a switchover using Device Manager, follow these steps:

Procedure

Step 1

Ensure that an HA switchover is possible by selecting Physical > Modules to verify the presence of multiple modules.

You see the screen shown in [Figure 5: Modules Screen Shows Current Supervisor](#), on page 238.

Figure 5: Modules Screen Shows Current Supervisor


Module	Name	Model	Status			StatusLastChangeTime	Power		
			Oper	Reset	ResetReasonDescription		Admin	Oper	Current
1	10 Gbps FC Module	DS-X9704	ok	<input type="checkbox"/>	Unknown	2006/02/22-11:21:31	on	ok	201.6W / 4.8A
4	1/2 Gbps FC Module	DS-X9016	ok	<input type="checkbox"/>	Unknown	2006/02/22-17:37:28	on	ok	210.0W / 5.0A
5	1/24 Gbps FC Module	DS-X9112	ok	<input type="checkbox"/>	Unknown reason	2006/02/22-11:56:56	on	ok	168.0W / 4.0A
7	Supervisor/Fabric-2	DS-X9530-SF2-K9	active	<input type="checkbox"/>	Reset Requested by CLI command reload	2006/02/22-11:13:47	on	ok	199.5W / 4.75A
8	Supervisor/Fabric-2	DS-X9530-SF2-K9	ha-standby	<input type="checkbox"/>	Unknown	2006/02/22-11:15:58	on	ok	199.5W / 4.75A
14	Fabric card module	DS-13SLT-FAB1	ok	<input type="checkbox"/>	Unknown	2006/02/22-11:13:56	on	ok	79.8W / 1.9A
15	Fabric card module	DS-13SLT-FAB1	ok	<input type="checkbox"/>	Module is powered down or power cycled	2006/02/22-17:43:56	on	ok	79.8W / 1.9A

7 row(s)

Apply Refresh Help Close

Step 2 In the main Device Manager screen, select **Admin > Reset Switch**.

Figure 6: Reset Switch Dialog Box



Step 3 Click **Switch to Standby**.

Copying Boot Variable Images to the Standby Supervisor Module

You can copy the boot variable images that are in the active supervisor module (but not in the standby supervisor module) to the standby supervisor module. Only those KICKSTART and SYSTEM boot variables that are set for the standby supervisor module can be copied. For module (line card) images, all boot variables are copied to the corresponding standby locations (bootflash: or slot0:) if not already present.

Displaying HA Status Information

The following conditions identify when automatic synchronization is possible:

- If the internal state of one supervisor module is Active with HA standby and the other supervisor module is HA standby, the switch is operationally HA and can do automatic synchronization.
- If the internal state of one of the supervisor modules is none, the switch cannot do automatic synchronization.

[Table 8: Redundancy States](#), on page 239 lists the possible values for the redundancy states.

Table 8: Redundancy States

State	Description
Not present	The supervisor module is not present or is not plugged into the chassis.
Initializing	The diagnostics have passed and the configuration is being downloaded.
Active	The active supervisor module and the switch is ready to be configured.
Standby	A switchover is possible.
Failed	<p>The switch detects a supervisor module failure on initialization and automatically attempts to power-cycle the module three (3) times. After the third attempt it continues to display a failed state.</p> <p>Note You should try to initialize the supervisor module until it comes up as HA-standby. This state is a temporary state.</p>

State	Description
Offline	The supervisor module is intentionally shut down for debugging purposes.
At BIOS	The switch has established connection with the supervisor and the supervisor module is performing diagnostics.
Unknown	The switch is in an invalid state. If it persists, call TAC.

[Table 9: Supervisor States](#) , on page 240 lists the possible values for the supervisor module states.

Table 9: Supervisor States

State	Description
Active	The active supervisor module in the switch is ready to be configured.
HA standby	A switchover is possible.
Offline	The switch is intentionally shut down for debugging purposes.
Unknown	The switch is in an invalid state and requires a support call to TAC.

[Table 10: Internal States](#) , on page 240 lists the possible values for the internal redundancy states.

Table 10: Internal States

State	Description
HA standby	The HA switchover mechanism in the standby supervisor module is enabled (see the Synchronizing Supervisor Modules , on page 238).
Active with no standby	A switchover is not possible.
Active with HA standby	The active supervisor module in the switch is ready to be configured. The standby supervisor module is in the HA-standby state.
Shutting down	The switch is being shut down.
HA switchover in progress	The switch is in the process of changing over to the HA switchover mechanism.
Offline	The switch is intentionally shut down for debugging purposes.
HA synchronization in progress	The standby supervisor module is in the process of synchronizing its state with the active supervisor modules.
Standby (failed)	The standby supervisor module is not functioning.
Active with failed standby	The active supervisor module and the second supervisor module is present but is not functioning.
Other	The switch is in a transient state. If it persists, call TAC.



CHAPTER 8

Configuring Trunking

- [Configuring Trunking, on page 241](#)

Configuring Trunking

This chapter describes how to configure trunking on E, F, N, and NP ports.

This chapter includes the following topics:

- [Information About Trunking, on page 241](#)
- [Guidelines and Limitations, on page 247](#)
- [Default Settings, on page 249](#)
- [Configuring Trunking, on page 250](#)


Information About Trunking

Trunking, also known as VSAN trunking, is a feature specific to switches in the Cisco MDS 9000 Family. Trunking enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link. Trunking is supported on E ports and F ports (See [Figure 7: Trunking E Ports, on page 241](#) and [Figure 8: Trunking F Ports, on page 242](#)).

Trunking E Ports

Trunking the E ports enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link, using enhanced ISL (EISL) frame format.

Figure 7: Trunking E Ports





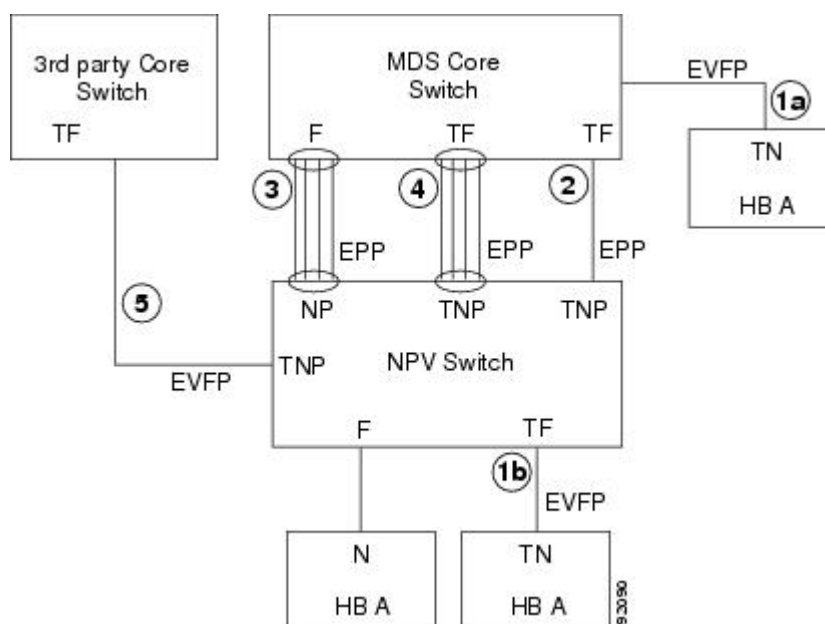
Note Trunking is not supported by internal ports on both the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter.

Trunking F Ports

Trunking F ports allows interconnected ports to transmit and receive tagged frames in more than one VSAN, over the same physical link.

Figure 8: Trunking F Ports, on page 242 represents the possible trunking scenarios in a SAN with MDS core switches, NPV switches, third-party core switches, and HBAs.

Figure 8: Trunking F Ports



Link Number	Link Description
1a and 1b	F port trunk with N port. ¹
2	F port trunk with NP port.
3	F PortChannel with NP port.
4	Trunked F PortChannel with NP port.
5	Trunking NP port with third-party core switch F port. ²

¹ These features are not supported currently.

² These features are not supported currently.

Key Concepts

The trunking feature includes the following key concepts:

- TE port—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.
- TF port—If trunk mode is enabled in an F port (see the link 2 in [Figure 8: Trunking F Ports, on page 242](#)) and that port becomes operational as a trunking F port, it is referred to as a TF port.
- TN port—If trunk mode is enabled (not currently supported) in an N port (see the link 1b in [Figure 8: Trunking F Ports, on page 242](#)) and that port becomes operational as a trunking N port, it is referred to as a TN port.
- TNP port—If trunk mode is enabled in an NP port (see the link 2 in [Figure 8: Trunking F Ports, on page 242](#)) and that port becomes operational as a trunking NP port, it is referred to as a TNP port.
- TF PortChannel—If trunk mode is enabled in an F PortChannel (see the link 4 in [Figure 8: Trunking F Ports, on page 242](#)) and that PortChannel becomes operational as a trunking F PortChannel, it is referred to as TF PortChannel. Cisco Port Trunking Protocol (PTP) is used to carry tagged frames
- TF-TN port link—A single link can be established to connect an F port to an HBA to carry tagged frames (see the link 1a and 1b in [Figure 8: Trunking F Ports, on page 242](#)) using Exchange Virtual Fabrics Protocol (EVFP). A server can reach multiple VSANs through a TF port without inter-VSAN routing (IVR).
- TF-TNP port link—A single link can be established to connect an TF port to an TNP port using the PTP protocol to carry tagged frames (see the link 2 in [Figure 8: Trunking F Ports, on page 242](#)). PTP is used because PTP also supports trunking PortChannels.



Note

The TF-TNP port link between a third-party NPV core and a Cisco NPV switch is established using the EVFP protocol.

- A Fibre Channel VSAN is called Virtual Fabric and uses a VF_ID in place of the VSAN ID. By default, the VF_ID is 1 for all ports. When an N port supports trunking, a pWWN is defined for each VSAN and called a logical pWWN. In the case of MDS core switches, the pWWNs for which the N port requests additional FC_IDs are called virtual pWWNs.

Trunking Protocols

The trunking protocol is important for trunking operations on the ports. The protocols enable the following activities:

- Dynamic negotiation of operational trunk mode.
- Selection of a common set of trunk-allowed VSANs.
- Detection of a VSAN mismatch across an ISL.

[Table 11: Supported Trunking Protocols, on page 244](#) specifies the protocols used for trunking and channeling.

Table 11: Supported Trunking Protocols

Trunk Link	Default
TE-TE port link	Cisco EPP (PTP)
TF-TN port link ³	FC-LS Rev 1.62 EVFP
TF-TNP port link	Cisco EPP (PTP)
E or F PortChannel	Cisco EPP (PCP)
TF Port Channel	Cisco EPP (PTP and PCP)
Third-party TF-TNP port link 1	FC-LS Rev 1.62 EVFP

³ These features are not currently supported.

By default, the trunking protocol is enabled on E ports and disabled on F ports. If the trunking protocol is disabled on a switch, no port on that switch can apply new trunk configurations. Existing trunk configurations are not affected. The TE port continues to function in trunk mode, but only supports traffic in VSANs that it negotiated with previously (when the trunking protocol was enabled). Also, other switches that are directly connected to this switch are similarly affected on the connected interfaces. In some cases, you may need to merge traffic from different port VSANs across a non-trunking ISL. If so, disable the trunking protocol.

**Note**

We recommend that both ends of a trunking link belong to the same port VSAN. On certain switches or fabric switches where the port VSANs are different, one end returns an error and the other end is not connected.

Trunk Modes

By default, trunk mode is enabled on all Fibre Channel interfaces (Mode: E, F, FL, Fx, ST, and SD) on non-NPV switches. On NPV switches, by default, trunk mode is disabled. You can configure trunk mode as on (enabled), off (disabled), or auto (automatic). The trunk mode configurations at the two ends of an ISL, between two switches, determine the trunking state of the link and the port modes at both ends. See below table.

Table 12: Trunk Mode Status Between Switches

Your Trunk Mode Configuration			Resulting State and Port Mode	
Port Type	Switch 1	Switch 2	Trunking State	Port Mode
E ports	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
	On	Auto or on	Trunking (EISL)	TE port
Port Type	Core Switch	NPV Switch	Trunking State	Link Mode

Your Trunk Mode Configuration			Resulting State and Port Mode	
Port Type	Switch 1	Switch 2	Trunking State	Port Mode
F and NP ports	On	Auto or on	Trunking	TF-TNP link
	Auto	On	Trunking	TF-TNP link
	Off	Auto, on, or off	No trunking	F-NP link

**Tip**

The preferred configuration on the Cisco MDS 9000 Family switches is one side of the trunk set to auto and the other side set to on.

**Note**

When connected to a third-party switch, the trunk mode configuration on E ports has no effect. The ISL is always in a trunking disabled state. In the case of F ports, if the third-party core switch ACC's physical FLOGI with the EVFP bit is configured, then EVFP protocol enables trunking on the link.

Trunk-Allowed VSAN Lists and VF_IDs

Each Fibre Channel interface has an associated trunk-allowed VSAN list. In TE-port mode, frames are transmitted and received in one or more VSANs specified in this list. By default, the VSAN range (1 through 4093) is included in the trunk-allowed list.

The common set of VSANs that are configured and active in the switch are included in the trunk-allowed VSAN list for an interface, and they are called *allowed-active* VSANs. The trunking protocol uses the list of allowed-active VSANs at the two ends of an ISL to determine the list of operational VSANs in which traffic is allowed.

Switch 1 (see [Figure 9: Default Allowed-Active VSAN Configuration, on page 246](#)) 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 (see [Figure 9: Default Allowed-Active VSAN Configuration, on page 246](#)).

For all F, N, and NP ports, the default VF_ID is 1 when there is no VF_ID configured. The trunk-allowed VF_ID list on a port is same as the list of trunk-allowed VSANs. VF_ID 4094 is called the control VF_ID and it is used to define the list of trunk-allowed VF-IDs when trunking is enabled on the link.

If F port trunking and channeling is enabled, or if **switchport trunk mode on** is configured in NPV mode for any interface, or if NP PortChannel is configured, the VSAN and VF-ID ranges available for the configuration are as described in [Table 13: VSAN and VF-ID Reservations, on page 245](#).

Table 13: VSAN and VF-ID Reservations

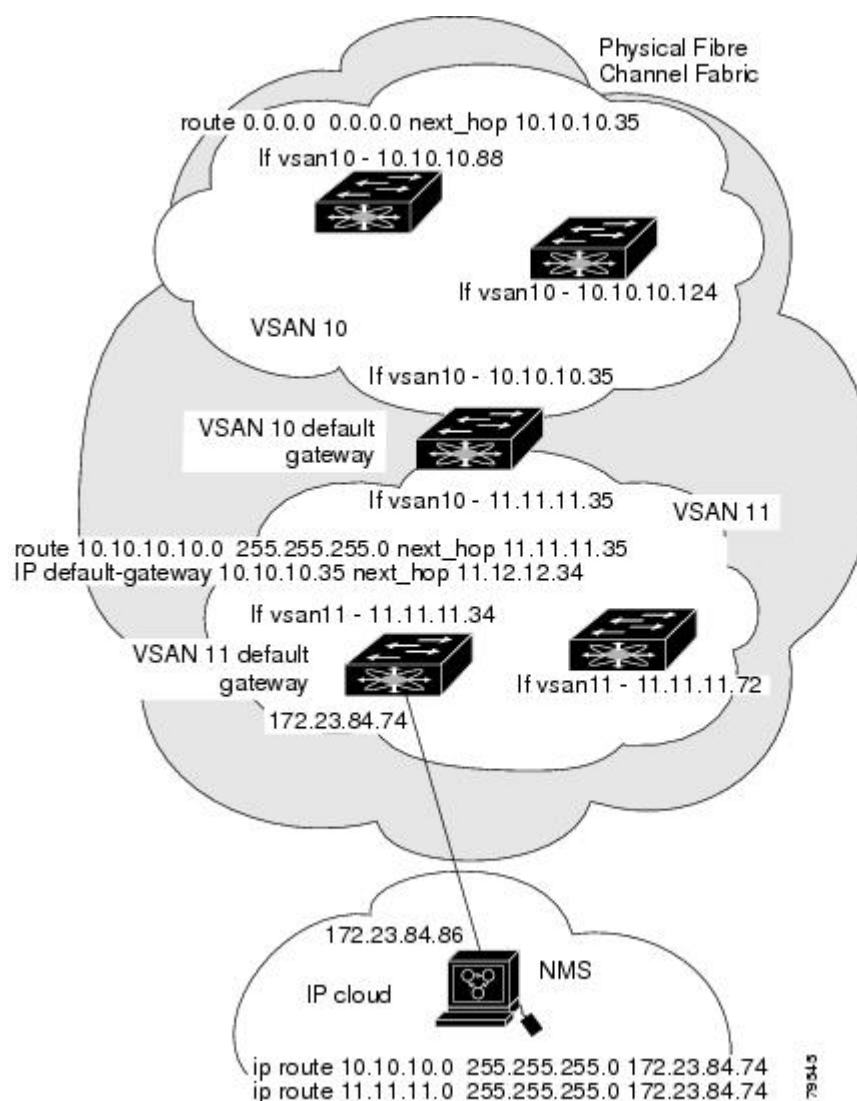
VSAN or VF-ID	Description
000h	Cannot be used as virtual fabric identifier.
001h(1) to EFFh(3839)	This VSAN range is available for user configuration.

VSAN or VF-ID	Description
F00h(3840) to FEEh(4078)	Reserved VSANs and they are not available for user configuration.
FEFh(4079)	EVFP isolated VSAN.
FF0h(4080) to FFEh(4094)	Used for vendor-specific VSANs.
FFFh	Cannot be used as virtual fabric identifier.



Note If the VF_ID of the F port and the N port do not match, then no tagged frames can be exchanged.

Figure 9: Default Allowed-Active VSAN Configuration




You can configure a select set of VSANs (from the allowed-active list) to control access to the VSANs specified in a trunking ISL.

Using [Figure 9: Default Allowed-Active VSAN Configuration, on page 246](#) as an example, you can configure the list of allowed VSANs on a per-interface basis (see [Figure 10: Operational and Allowed VSAN Configuration, on page 247](#)). 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.

Figure 10: Operational and Allowed VSAN Configuration



Guidelines and Limitations

Trunking has the following configuration guidelines and limitations:

General Guidelines and Limitations

The trunking feature has the following general configuration guidelines and limitations:

- F ports support trunking in Fx mode.
- The trunk-allowed VSANs configured for TE, TF, and TNP links 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.
- Trunking F ports and trunking F PortChannels are not supported on the following hardware:

- 91x4 switches, if NPIV is enabled and used as the NPIV core switch.
- Generation 1 2-Gbps Fibre Channel switching modules.
- On core switches, the FC-SP authentication will be supported only for the physical FLOGI from the physical pWWN.
- No FC-SP authentication is supported by the NPV switch on the server F ports.
- MDS does not enforce the uniqueness of logical pWWNs across VSANs.
- DPVM is not supported on trunked F port logins.
- The DPVM feature is limited to the control of the port VSAN, since the EVFP protocol does not allow changing the VSAN on which a logical pWWN has done FLOGI.
- The port security configuration will be applied to both the first physical FLOGI and the per VSAN FLOGIs.
- Trunking is not supported on F ports that have FlexAttach enabled.
- On MDS 91x4 core switches, hard zoning can be done only on F ports that are doing either NPIV or trunking. However, in NPV mode, this restriction does not apply since zoning is enforced on the core F port.

Upgrade and Downgrade Limitations

The trunking and channeling feature includes the following upgrade and downgrade limitations:

- When F port trunking or channeling is configured on a link, the switch cannot be downgraded to Cisco MDS SAN-OS Release 3.x and NX-OS Release 4.1(1b), or earlier.
- If you are upgrading from a SAN-OS Release 3.x to NX-OS Release 5.0(1), and you have not created VSAN 4079, the NX-OS software will automatically create VSAN 4079 and reserve it for EVFP use.
 - If you have created VSAN 4079, the upgrade to NX-OS Release 5.0(1) will have no affect on VSAN 4079.
 - If you downgrade after NX-OS Release 5.0(1), the VSAN will no longer be reserved for EVFP use.

Difference Between TE Ports and TF-TNP Ports

In case of TE ports, the VSAN will be in initializing state when VSAN is coming up on that interface and when peers are in negotiating phase. Once the handshake is done, VSAN will be moved to up state in the successful case, and isolated state in the case of failure. Device Manager will show the port status as amber during initializing state and it will be green once VSANs are up.

In case of TF ports, after the handshake, one of the allowed VSANs will be moved to the up state. All other VSANs will be in initializing state even though the handshake with the peer is completed and successful. Each VSAN will be moved from initializing state to up state when a server or target logs in through the trunked F or NP ports in the corresponding VSAN.




Note

In case of TF or TNP ports, the Device Manager will show the port status as amber even after port is up and there is no failure. It will be changed to green once all the VSAN has successful logins.

Trunking Misconfiguration Examples

If you do not configure the VSANs correctly, issues with the connection may occur. For example, if you merge the traffic in two VSANs, both VSANs will be mismatched. The trunking protocol validates the VSAN interfaces at both ends of a link to avoid merging VSANs (see [Figure 11: VSAN Mismatch, on page 249](#)).

Figure 11: VSAN Mismatch



The trunking protocol detects potential VSAN merging and isolates the ports involved (see [Figure 11: VSAN Mismatch, on page 249](#)).

The trunking protocol cannot detect merging of VSANs when a third-party switch is placed in between two Cisco MDS 9000 Family switches (see [Figure 12: Third-Party Switch VSAN Mismatch, on page 249](#)).

Figure 12: Third-Party Switch VSAN Mismatch



VSAN 2 and VSAN 3 are effectively merged with overlapping entries in the name server and the zone applications. Cisco DCNM-SAN helps detect such topologies.

Default Settings

[Table 14: Default Trunk Configuration Parameters , on page 249](#) lists the default settings for trunking parameters.

Table 14: Default Trunk Configuration Parameters

Parameters	Default
Switch port trunk mode	ON on non-NPV and MDS core switches. OFF on NPV switches.
Allowed VSAN list	1 to 4093 user-defined VSAN IDs.
Allowed VF-ID list	1 to 4093 user-defined VF-IDs.
Trunking protocol on E ports	Enabled.
Trunking protocol on F ports	Disabled.

Configuring Trunking

This section includes the following topics:

Enabling the Cisco Trunking and Channeling Protocols

This section describes how to enable the required trunking and channeling protocols.

Prerequisites

To avoid inconsistent configurations, disable all ports with a **shutdown** command before enabling or disabling the trunking protocols.

Detailed Steps

Procedure

	Command or Action	Purpose
Step 1	To enable or disable the Cisco trunking and channeling protocol, follow these steps:	

Enabling the F Port Trunking and Channeling Protocol

To enable or disable the F port trunking and channeling protocols using DCNM-SAN, follow these steps:

Procedure

-
- Step 1** From the Physical Attributes panel, expand Switches. Select **FC Services**, and then select F_Port_Channel/Trunk.
- You see the list of switches in the Fabric with F port trunking and channeling enabled.
- Step 2** From the command column, select **enable** or **disable** or **no selection**.
-

Configuring Trunk Mode

To configure trunk mode using DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Expand **FC Interfaces**, and then select **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.
- Step 3** Make changes to the Admin and Allowed VSANs values.
- Step 4** Click the **Trunk Failures** tab to check if a link did not come up.
- You see the reason listed in the FailureCause column.

Step 5 Click the **Apply Changes** icon.



CHAPTER 9

Configuring PortChannels

- [Configuring PortChannels, on page 253](#)

Configuring PortChannels


This chapter describes how to configure PortChannels.

Information About PortChannels

PortChannels Overview

PortChannels refer to the aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy (See [Figure 13: PortChannel Flexibility, on page 253](#)). PortChannels can connect to interfaces across switching modules, so a failure of a switching module cannot bring down the PortChannel link.

Figure 13: PortChannel Flexibility



PortChannels on Cisco MDS 9000 Family switches allow flexibility in configuration. This illustrates three possible PortChannel configurations:

- PortChannel A aggregates two links on two interfaces on the same switching module at each end of a connection.

- PortChannel B also aggregates two links, but each link is connected to a different switching module. If the switching module goes down, traffic is not affected.
- PortChannel C aggregates three links. Two links are on the same switching module at each end, while one is connected to a different switching module on switch 2.

E PortChannels

An E PortChannel refers to the aggregation of multiple E ports into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy. PortChannels can connect to interfaces across switching modules, so a failure of a switching module cannot bring down the PortChannel link.

A PortChannel has the following features and restrictions:

- Provides a point-to-point connection over ISL (E ports) or EISL (TE ports). Multiple links can be combined into a PortChannel.
- Increases the aggregate bandwidth on an ISL by distributing traffic among all functional links in the channel.
- Load balances across multiple links and maintains optimum bandwidth utilization. Load balancing is based on the source ID, destination ID, and exchange ID (OX ID).
- Provides high availability on an ISL. If one link fails, traffic previously carried on this link is switched to the remaining links. If a link goes down in a PortChannel, the upper protocol is not aware of it. To the upper protocol, the link is still there, although the bandwidth is diminished. The routing tables are not affected by link failure. PortChannels may contain up to 16 physical links and may span multiple modules for added high availability.



Note See the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide* for information about failover scenarios for PortChannels and FSPF links.

F and TF PortChannels

An F PortChannel is also a logical interface that combines a set of F ports connected to the same Fibre Channel node and operates as one link between the F ports and the NP ports. The F PortChannels support bandwidth utilization and availability like the E PortChannels. F PortChannels are mainly used to connect MDS core and NPV switches to provide optimal bandwidth utilization and transparent failover between the uplinks of a VSAN.

An F PortChannel trunk combines the functionality and advantages of a TF port and an F PortChannel. This logical link uses the Cisco PTP and PCP protocols over Cisco EPP (ELS).



Note If a Cisco MDS 9124 or 9134 switch is used as a core switch, only a nontrunking F PortChannel is supported. Trunking is not supported on this platform when NPIV enabled.

PortChanneling and Trunking

Trunking is a commonly used storage industry term. However, the Cisco NX-OS software and switches in the Cisco MDS 9000 Family implement trunking and PortChanneling as follows:

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

- Trunking enables a link transmitting frames in the EISL format to carry (trunk) multiple VSAN traffic. For example, when trunking is operational on an E port, that E port becomes a TE port. A TE port is specific to switches in the Cisco MDS 9000 Family. An industry standard E port can link to other vendor switches and is referred to as a nontrunking interface (see [Figure 14: Trunking Only, on page 255](#) and [Figure 15: PortChanneling and Trunking, on page 255](#)). See Chapter 8, “Configuring Trunking,” for information on trunked interfaces.

Figure 14: Trunking Only





Figure 15: PortChanneling and Trunking



PortChanneling and trunking are used separately across an ISL.

- PortChanneling—Interfaces can be channeled between the following sets of ports:
 - E ports and TE ports
 - F ports and NP ports
 - TF ports and TNP ports
- Trunking—Trunking permits carrying traffic on multiple VSANs between switches. See the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.
- Both PortChanneling and trunking can be used between TE ports over EISLs.

Load Balancing

Two methods support the load-balancing functionality:

- Flow based—All frames between source and destination follow the same links for a given flow. That is, whichever link is selected for the first exchange of the flow is used for all subsequent exchanges.
- Exchange based—The first frame in an exchange picks a link and subsequent frames in the exchange follow the same link. However, subsequent exchanges can use a different link. This provides more granular load balancing while preserving the order of frames for each exchange.

The following image illustrates how source ID 1 (SID1) and destination ID 1 (DID1) 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 16: SID1 and DID1-Based Load Balancing


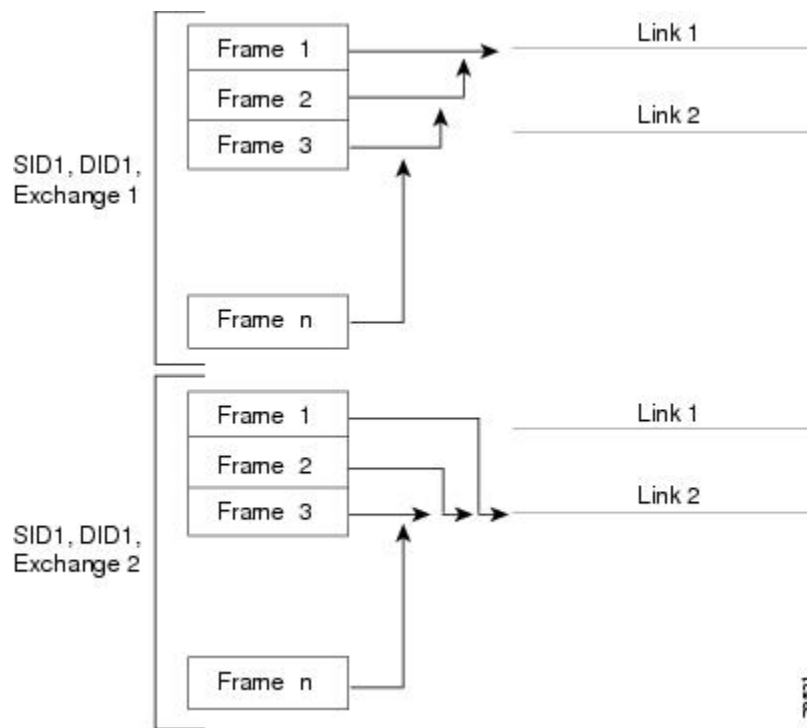


Figure 17: SID1, DID1, and Exchange-Based Load Balancing , on page 257 illustrates how exchange-based load balancing works. When the first frame in an exchange is received for forwarding on an interface, link 1 is chosen by a hash algorithm. All remaining frames in that particular exchange are sent on 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 17: SID1, DID1, and Exchange-Based Load Balancing



For more information on configuring load balancing and in-order delivery features, see the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.

PortChannel Modes

You can configure each PortChannel with a channel group mode parameter to determine the PortChannel protocol behavior for all member ports in this channel group. The possible values for a channel group mode are as follows:

- **ON (default)**—The member ports only operate as part of a PortChannel or remain inactive. In this mode, the PortChannel protocol is not initiated. However, if a PortChannel protocol frame is received from a peer port, the software indicates its nonnegotiable status. This mode is backward compatible with the existing implementation of PortChannels in releases prior to Release 2.0(1b), where the channel group mode is implicitly assumed to be ON. In Cisco MDS SAN-OS Releases 1.3 and earlier, the only available PortChannel mode was the ON mode. PortChannels configured in the ON mode require you to explicitly enable and disable the PortChannel member ports at either end if you add or remove ports from the PortChannel configuration. You must physically verify that the local and remote ports are connected to each other.
- **ACTIVE**—The member ports initiate PortChannel protocol negotiation with the peer port(s) regardless of the channel group mode of the peer port. If the peer port, while configured in a channel group, does not support the PortChannel protocol, or responds with a nonnegotiable status, it will default to the ON mode behavior. The ACTIVE PortChannel mode allows automatic recovery without explicitly enabling and disabling the PortChannel member ports at either end.

Table 15: Channel Group Configuration Differences, on page 258 compares ON and ACTIVE modes.

Table 15: Channel Group Configuration Differences

ON Mode	ACTIVE Mode
No protocol is exchanged.	A PortChannel protocol negotiation is performed with the peer ports.
Moves interfaces to the suspended state if its operational values are incompatible with the PortChannel.	Moves interfaces to the isolated state if its operational values are incompatible with the PortChannel.
When you add or modify a PortChannel member port configuration, you must explicitly disable (shut) and enable (no shut) the PortChannel member ports at either end.	When you add or modify a PortChannel interface, the PortChannel automatically recovers.
Port initialization is not synchronized.	There is synchronized startup of all ports in a channel across peer switches.
All misconfigurations are not detected as no protocol is exchanged.	Consistently detect misconfigurations using a PortChannel protocol.
Transitions misconfigured ports to the suspended state. You must explicitly disable (shut) and enable (no shut) the member ports at either end.	Transitions misconfigured ports to the isolated state to correct the misconfiguration. Once you correct the misconfiguration, the protocol ensures automatic recovery.
This is the default mode.	You must explicitly configure this mode.

PortChannel Deletion

When you delete the PortChannel, the corresponding channel membership is also deleted. All interfaces in the deleted PortChannel convert to individual physical links. After the PortChannel is removed, regardless of the mode used (ACTIVE and ON), the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the [“Graceful Shutdown” section on page 11-9](#)).

If you delete the PortChannel for one port, then the individual ports within the deleted PortChannel retain the compatibility parameter settings (speed, mode, port VSAN, allowed VSAN, and port security). You can explicitly change those settings as required.

- If you use the default ON mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, then the PortChannel ports automatically recover from the deletion.

Interfaces in a PortChannel

You can add or remove a physical interface (or a range of interfaces) to an existing PortChannel. The compatible parameters on the configuration are mapped to the PortChannel. Adding an interface to a PortChannel increases the channel size and bandwidth of the PortChannel. Removing an interface from a PortChannel decreases the channel size and bandwidth of the PortChannel.

This section describes interface configuration for a PortChannel and includes the following topics:



Note For information about PortChannel support on Generation 2 switching modules, see the [“PortChannel Limitations” section on page 12-23](#).

Interface Addition to a PortChannel

You can add a physical interface (or a range of interfaces) to an existing PortChannel. The compatible parameters on the configuration are mapped to the PortChannel. Adding an interface to a PortChannel increases the channel size and bandwidth of the PortChannel.

A port can be configured as a member of a static PortChannel only if the following configurations are the same in the port and the PortChannel:

- Speed
- Mode
- Rate mode
- Port VSAN
- Trunking mode
- Allowed VSAN list or VF-ID list

After the members are added, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the [Generation 1 PortChannel Limitations, on page 264](#) and [“Graceful Shutdown” section on page 11-9](#)).

Compatibility Check

A compatibility check ensures that the same parameter settings are used in all physical ports in the channel. Otherwise, they cannot become part of a PortChannel. The compatibility check is performed before a port is added to the PortChannel.

The check ensures that the following parameters and settings match at both ends of a PortChannel:

- Capability parameters (type of interface, Gigabit Ethernet at both ends, or Fibre Channel at both ends).
- Administrative compatibility parameters (speed, mode, rate mode, port VSAN, allowed VSAN list, and port security).



Note Ports in shared rate mode cannot form a PortChannel or a trunking PortChannel.

- Operational parameters (remote switch WWN and trunking mode).

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.

Forcing an Interface Addition

You can force the port configuration to be overwritten by the PortChannel. In this case, the interface is added to a PortChannel.

- If you use the default ON mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, then the PortChannel ports automatically recover from the addition.



Note

When PortChannels are created from within an interface, the **force** option cannot be used.

After the members are forcefully added, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the [“Graceful Shutdown” section on page 11-9](#) and [Generation 1 PortChannel Limitations, on page 264](#)).

Interface Deletion from a PortChannel

When a physical interface is deleted from the PortChannel, the channel membership is automatically updated. If the deleted interface is the last operational interface, then the PortChannel status is changed to a down state. Deleting an interface from a PortChannel decreases the channel size and bandwidth of the PortChannel.

- If you use the default ON mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, then the PortChannel ports automatically recover from the deletion.

After the members are deleted, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the [Generation 1 PortChannel Limitations, on page 264](#) and [“Graceful Shutdown” section on page 11-9](#)).

PortChannel Protocols

In earlier Cisco SAN-OS releases, PortChannels required additional administrative tasks to support synchronization. The Cisco NX-OS software provides robust error detection and synchronization capabilities. You can manually configure channel groups or they can be automatically created. In both cases, the channel groups have the same capability and configurational parameters. Any change in configuration applied to the associated PortChannel interface is propagated to all members of the channel group.

A protocol to exchange PortChannel configurations is available in all Cisco MDS switches. This addition simplifies PortChannel management with incompatible ISLs. An additional autcreation mode enables ISLs with compatible parameters to automatically form channel groups without manual intervention.

The PortChannel protocol is enabled by default.

The PortChannel protocol expands the PortChannel functional model in Cisco MDS switches. It uses the exchange peer parameters (EPP) services to communicate across peer ports in an ISL. Each switch uses the information received from the peer ports along with its local configuration and operational values to decide if it should be part of a PortChannel. The protocol ensures that a set of ports are eligible to be part of the same PortChannel. They are only eligible to be part of the same PortChannel if all the ports have a compatible partner.

The PortChannel protocol uses two subprotocols:

- Bringup protocol—Automatically detects misconfigurations so you can correct them. This protocol synchronizes the PortChannel at both ends so that all frames for a given flow (as identified by the source

FC ID, destination FC ID and OX_ID) are carried over the same physical link in both directions. This helps make applications such as write acceleration, work for PortChannels over FCIP links.

- Autocreation protocol—Automatically aggregates compatible ports into a PortChannel.

This section describes how to configure the PortChannel protocol and includes the following sections:


Channel Group Creation



Note Channel groups are not supported on internal ports in the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeSystem.

Assuming link A1-B1 comes up first (see [Figure 18: Autocreating Channel Groups, on page 261](#)), that link is operational as an individual link. When the next link comes up, for example, A2-B2, the PortChannel protocol identifies if this link is compatible with link A1-B1 and automatically creates channel groups 10 and 20 in the respective switches. If link A3-B3 can join the channel groups (the PortChannels), the respective ports have compatible configurations. If link A4-B4 operates as an individual link, it is because of the incompatible configuration of the two end ports with the other member ports in this channel group.

Figure 18: Autocreating Channel Groups



The channel group numbers are selected dynamically, and as such, the administrative configuration of the ports forming the channel group at either end are applicable to the newly created channel group. The channel group number being chosen dynamically may be different across reboots for the same set of PortChannels based on the order of ports that are initialized in the switch.

[Table 16: Channel Group Configuration Differences , on page 261](#) identifies the differences between user-configured and auto-configured channel groups.

Table 16: 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.

User-Configured Channel Group	Autocreated Channel Group
You can form the PortChannel with a subset of the ports in the channel group. Incompatible ports remain in a suspended or isolated state depending on the ON or ACTIVE mode configuration.	All ports included in the channel group participate in the PortChannel—no member port becomes isolated or suspended; instead, the member port is removed from the channel group when the link is found to be incompatible.
Any administrative configuration made to the PortChannel is applied to all ports in the channel group, and you can save the configuration for the PortChannel interface.	Any administrative configuration made to the PortChannel is applied to all ports in the channel group, but the configurations are saved for the member ports; no configuration is saved for the PortChannel interface. You can explicitly convert this channel group, if required.
You can remove any channel group and add members to a channel group.	You cannot remove a channel group, or add/remove any of its members. The channel group is removed when no member ports exist.

**Note**

Autocreation is not supported as of MDS NX-OS Release 4.1(1b) and later.

Autocreation

The autocreation protocol has the following functionality:

- A port is not allowed to be configured as part of a PortChannel when the autocreation feature is enabled. These two configurations are mutually exclusive.
- Autocreation must be enabled in both the local and peer ports to negotiate a PortChannel.
- Aggregation occurs in one of two ways:
 - A port is aggregated into a compatible autocreated PortChannel.
 - A port is aggregated with another compatible port to form a new PortChannel.
- Newly created PortChannels are allocated from the maximum possible PortChannel (128 for Generation 1 or a combination of Generation 1 and Generation 2 switches, or 256 for Generation 2 switches) in a decreasing order based on availability. If all 128 (or 256) numbers are used up, aggregation is not allowed.
- You cannot change the membership or delete an autocreated PortChannel.
- When you disable autocreation, all member ports are removed from the autocreated PortChannel.
- Once the last member is removed from an autocreated PortChannel, the channel is automatically deleted and the number is released for reuse.
- An autocreated PortChannel is not persistent through a reboot. An autocreated PortChannel can be manually configured to appear the same as a persistent PortChannel. Once the PortChannel is made persistent, the autocreation feature is disabled in all member ports.
- 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 MDS 9000 Family, we recommend that you retain at least one interconnected port between the switches without any autocreation configuration. If all ports between two switches are configured with the autocreation feature at the same time, you may face a possible traffic disruption between these two switches as the ports are automatically disabled and reenabled when ports are added to an autocreated PortChannel.

Manually Configured Channel Groups

A user-configured channel group cannot be converted to an autocreated channel group. However, you can convert an autocreated channel group to a manual channel group. Once performed, this task is irreversible. The channel group number does not change, but the member ports operate according to the properties of the manually configured channel group, and the autocreation of channel group is implicitly disabled for all member ports.



Tip If you enable persistence, be sure to enable it at both ends of the PortChannel.

Prerequisites for PortChannels

Before configuring a PortChannel, consider the following guidelines:

- Configure the PortChannel across switching modules to implement redundancy on switching module reboots or upgrades.
- Ensure that one PortChannel is not connected to different sets of switches. PortChannels require point-to-point connections between the same set of switches.



Note On switches with Generation 1 switching modules, or a combination of Generation 1 and Generation 2 switching modules, you can configure a maximum of 128 PortChannels. On switches with only Generation 2 switching modules, or Generation 2 and Generation 3 switching modules, you can configure a maximum of 256 PortChannels.

If you misconfigure PortChannels, you may receive a misconfiguration message. If you receive this message, the PortChannel's physical links are disabled because an error has been detected.

A PortChannel error is detected if the following requirements are not met:

- Each switch on either side of a PortChannel must be connected to the same number of interfaces.
- Each interface must be connected to a corresponding interface on the other side (see the *Valid and Invalid PortChannel Examples* section for an example of an invalid configuration).
- Links in a PortChannel cannot be changed after the PortChannel is configured. If you change the links after the PortChannel is configured, be sure to reconnect the links to interfaces within the PortChannel and reenabling the links.

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

Enter the **show interface** command for that interface to verify that the PortChannel is functioning as required.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

General Guidelines and Limitations

Cisco MDS 9000 Family switches support the following number of PortChannels per switch:

- Switches with only Generation 1 switching modules do not support F and TF PortChannels.
- Switches with Generation 1 switching modules, or a combination of Generation 1 and Generation 2 switching modules, support a maximum of 128 PortChannels. Only Generation 2 ports can be included in the PortChannels.
- Switches with only Generation 2 switching modules or Generation 2 and Generation 3 modules support a maximum of 256 PortChannels with 16 interfaces per PortChannel.
- A PortChannel number refers to the unique identifier for each channel group. This number ranges from of 1 to 256.

Generation 1 PortChannel Limitations

This section includes the restrictions on creation and addition of PortChannel members to a PortChannel on Generation 1 hardware:

- The 32-port 2-Gbps or 1-Gbps switching module.
- The MDS 9140 and 9120 switches.

When configuring the host-optimized ports on Generation 1 hardware, the following PortChannel guidelines apply:

- If you execute the **write erase** command on a 32-port switching module, and then copy a saved configuration to the switch from a text file that contains the **no system default switchport shutdown** command, you need to copy the text file to the switch again for the E ports to come up without manual configuration.
- Any (or all) full line rate port(s) in the Cisco MDS 9100 Series can be included in a PortChannel.
- The host-optimized ports in the Cisco MDS 9100 Series are subject to the same PortChannel rules as 32-port switching modules; only the first port of each group of 4 ports is included in a PortChannel.
 - You can configure only the first port in each 4-port group as an E port (for example, the first port in ports 1–4, the fifth port in ports 5–8, and so on). If the first port in the group is configured as a PortChannel, the other three ports in each group (ports 2–4, 6–8, and so on) are not usable and remain in the shutdown state.
 - If any of the other three ports are configured in a no shutdown state, you cannot configure the first port to be a PortChannel. The other three ports continue to remain in a no shutdown state.

F and TF PortChannel Limitations

The following guidelines and restrictions are applicable for F and TF PortChannels:

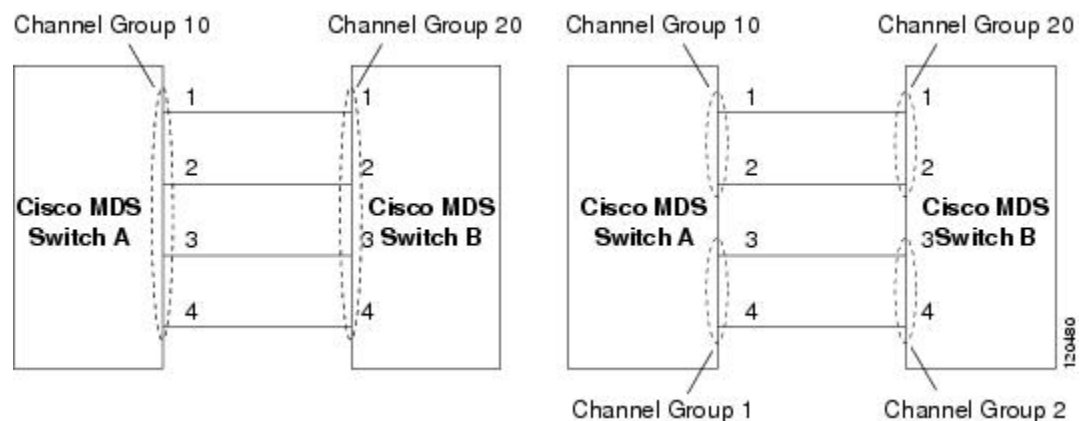
- The ports must be in F mode.
- Automatic creation is not supported.
- The PortChannel interface must be in ACTIVE mode when multiple FCIP interfaces are grouped with WA.
- ON mode is not supported. Only ACTIVE-ACTIVE mode is supported. By default, the mode is ACTIVE on the NPV switches.

- Devices logged in through F PortChannel on an MDS switch are not supported in IVR non-NAT configuration. The devices are supported only in IVR NAT configuration.
- Port security rules are enforced only on physical pWWNs at the single link level.
- FC-SP authenticates only the first physical FLOGI of every PortChannel member.
- Since the FLOGI payload carries only the VF bits to trigger the use of a protocol after the FLOGI exchange, those bits will be overridden. In the case of the NPV switches, the core has a Cisco WWN and will try to initiate the PCP protocol.
- The name server registration of the N ports logging in through an F PortChannel will use the fWWN of the PortChannel interface.
- DPVM configuration is not supported.
- The PortChannel port VSAN cannot be configured using DPVM.
- The Dynamic Port VSAN Management (DPVM) database will be queried only for the first physical FLOGI of each member, so that the port VSAN can be configured automatically.
- DPVM does not bind FC_IDs to VSANs, but pWWNs to VSANs. It will be queried only for the physical FLOGI.

Valid and Invalid PortChannel Examples


PortChannels are created with default values. You can change the default configuration just like any other physical interface. [Figure 19: Valid PortChannel Configurations, on page 265](#) provides examples of valid PortChannel configurations.

Figure 19: Valid PortChannel Configurations



[Figure 20: Misconfigured Configurations, on page 266](#) provides examples of invalid configurations. Assuming that the links are brought up in the 1, 2, 3, 4 sequence, links 3 and 4 will be operationally down as the fabric is misconfigured.

Figure 20: Misconfigured Configurations



130488

Default Settings

Table 17: Default PortChannel Parameters , on page 266 lists the default settings for PortChannels.

Table 17: Default PortChannel Parameters

Parameters	Default
PortChannels	FSPF is enabled by default.
Create PortChannel	Administratively up.
Default PortChannel mode	ON mode on non-NPV and NPV core switches. ACTIVE mode on NPV switches.
Autocreation	Disabled.

Configuring PortChannels

Configuring PortChannels Using the Wizard

To create a PortChannel using the PortChannel Wizard in DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Click the PortChannel Wizard icon in the toolbar.
You see the first PortChannel Wizard screen.
- Step 2** Select a switch pair.
- Step 3** Click **Next**.
- Step 4** Select the ISLs.
- Step 5** (Optional) Check the Dynamically form Port Channel Group from selected ISLs check box if you want to dynamically create the PortChannel and make the ISL properties identical for the Admin, Trunk, Speed, and VSAN attributes.
- Step 6** Click **Next**.
- Step 7** If you chose to dynamically form a PortChannel from selected ISLs, you see the final PortChannel Wizard screen. Set the VSAN List, Trunk Mode, and Speed and proceed to Step 11.
- Step 8** If you did not choose to dynamically form a PortChannel, you see the third PortChannel Wizard dialog box.
- Note** Dynamic VSAN creation is not supported on NPV switches.
- 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 the screen:
- VSAN List—Provides a list of VSANs to which the ISLs belong.
 - Trunk Mode—Enables trunking on the links in the PortChannel. Select **trunking** if your link is between TE ports. Select **nontrunking** if your link is between E ports. Select **auto** if you are not sure.
 - Force Admin, Trunk, Speed, and VSAN attributes to be identical—Ensures that the same parameter settings are used in all physical ports in the channel. If these settings are not identical, the ports cannot become part of the PortChannel.
 - Speed—The port speed values are auto, 1Gb, 2Gb, 4Gb, 8Gb, autoMax2G, and autoMax4G.
- Step 11** Click **OK**.
The PortChannel is created. Note that it may take a few minutes before the new PortChannel is visible in the Fabric pane.
-

Configuring the PortChannel Mode

To configure ACTIVE mode using DCNM-SAN, follow these steps:



Note

An F PortChannel is supported only on ACTIVE mode.

Procedure

-
- Step 1** Expand ISLs and then select Port Channels in the Physical Attributes pane.

You see the PortChannels configured in the Information pane.

- Step 2** Click the Protocols tab, and then from the Mode drop-down menu, select the appropriate mode for the PortChannel.
- Step 3** Click the Apply Changes icon to save any modifications.
-

Deleting PortChannels

To delete a PortChannel using the PortChannel Wizard in DCNM-SAN, follow these steps:

Procedure

- Step 1** Click the PortChannel Wizard icon in the toolbar.
You see the first PortChannel Wizard screen.
- Step 2** Select the existing PortChannel that you want to delete and click Next. You see a list of the ISLs currently associated with this PortChannel.
- Step 3** Click Next. You see an editable list of associated ISLs and available ISLs for this PortChannel.
- Step 4** Click each associated ISL and click the left arrow to remove all ISLs from the PortChannel.
- Step 5** Check the Delete Port Channel If Empty check box to delete this PortChannel.
- Step 6** Click Finish to save any modifications or click Cancel to discard any changes.
-

Adding an Interface to a PortChannel



Note

To add a range of ports to a PortChannel, follow these steps: By default, the CLI adds an interface normally to a PortChannel, while DCNM-SAN adds the interface by force, unless specified explicitly.

To add an interface or range of interfaces to a PortChannel using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand ISLs and then select Port Channels in the Physical Attributes pane.
You see the PortChannels in the Information pane.
- Step 2** Click the Channels tab and find the switch and PortChannel that you want to edit.
- Step 3** Set Members Admin to the interface or list of interfaces that you want to add to the PortChannel.
- Step 4** Click the Apply Changes icon to save any modifications or click Undo Changes to discard any changes.
-

Forcing an Interface Addition

To force the addition of a port to a PortChannel using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand ISLs and then select Port Channels in the Physical Attributes pane. You see the PortChannels in the Information pane.
 - Step 2** Click the Channels tab and find the switch and PortChannel that you want to edit.
 - Step 3** Set Members Admin to the interface or list of interfaces that you want to add to the PortChannel.
 - Step 4** Check the Force check box to force this interface addition.
 - Step 5** Click the Apply Changes icon to save any modifications.
-

Deleting an Interface from a PortChannel

To delete a physical interface (or a range of physical interfaces) from a PortChannel using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand ISLs and then select Port Channels in the Physical Attributes pane.
You see the PortChannels in the Information pane.
 - Step 2** Click the Channels tab and find the switch and PortChannel that you want to edit.
 - Step 3** Remove the interface or list of interfaces you want deleted in the Members the Admin column.
 - Step 4** Click the Apply Changes icon to save any modifications.
-

Converting to Manually Configured Channel Groups

You can convert autocreated channel group to a user-configured channel group using the **port-channel channel-group-number persistent** EXEC command. If the PortChannel does not exist, this command is not executed.

To convert an autocreated channel group to a user-configured channel group using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand ISLs and then select Port Channels in the Physical Attributes pane.
- Step 2** Click the Protocol tab.
You see the switch protocols.

- Step 3** Check the Persist check box for each channel that you want to convert to a manually configured channel group.
- Step 4** Click the Apply Changes icon to save any modifications.
-



CHAPTER 10

Configuring N Port Virtualization

- [Configuring N Port Virtualization, on page 271](#)

Configuring N Port Virtualization

This chapter describes how to configure N port virtualization.

Information About N Port Virtualization

NPV Overview

N port virtualization (NPV) reduces the number of Fibre Channel domain IDs in SANs. Switches operating in the NPV mode do not join a fabric. They pass traffic between NPV core switch links and end devices, which eliminates the domain IDs for these edge switches.

NPV is supported by the following Cisco MDS 9000 switches and Cisco Nexus 5000 Series switches only:

- Cisco MDS 9124 Multilayer Fabric Switch
- Cisco MDS 9134 Fabric Switch
- Cisco MDS 9148 Multilayer Fabric Switch
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter
- Cisco Nexus 5000 Series switches




Note NPV is available on these switches only while in NPV mode; if in switch mode, NPV is not available.

N Port Identifier Virtualization

N port identifier virtualization (NPIV) provides a means to assign multiple FC IDs to a single N port. This feature allows multiple applications on the N port to use different identifiers and allows access control, zoning, and port security to be implemented at the application level.

[Figure 21: NPIV Example, on page 272](#) shows an example application using NPIV.

Figure 21: NPIV Example



You must globally enable NPIV for all VSANs on the MDS switch to allow the NPIV-enabled applications to use multiple N port identifiers.

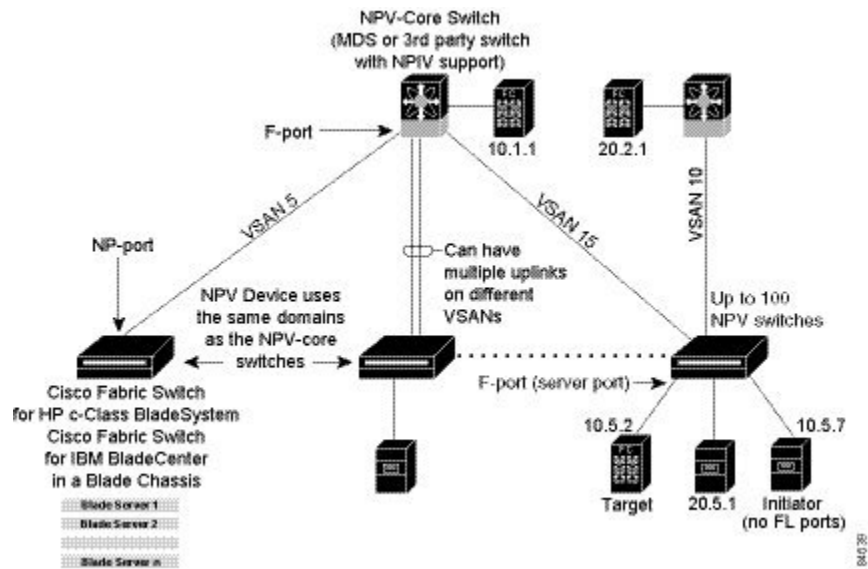
**Note**

All of the N port identifiers are allocated in the same VSAN.

N Port Virtualization


Typically, Fibre Channel networks are deployed using a core-edge model with a large number of fabric switches connected to edge devices. Such a model is cost-effective because the per port cost for director class switches is much higher than that of fabric switches. However, as the number of ports in the fabric increases, the number of switches deployed also increases, and you can end up with a significant increase in the number of domain IDs (the maximum number supported is 239). This challenge becomes even more difficult when additional blade chassis are deployed in Fibre Channel networks.

NPV addresses the increase in the number of domain IDs needed to deploy a large number of the ports by making a fabric or blade switch appear as a host to the core Fibre Channel switch, and as a Fibre Channel switch to the servers in the fabric or blade switch. NPV aggregates multiple locally connected N ports into one or more external NP links, which shares the domain ID of the NPV core switch among multiple NPV switches. NPV also allows multiple devices to attach to same port on the NPV core switch, which reduces the need for more ports on the core.



While NPV is similar to N port identifier virtualization (NPIV), it does not offer exactly the same functionality. NPIV provides a means to assign multiple FC IDs to a single N port, and allows multiple applications on the N port to use different identifiers. NPIV also allows access control, zoning, and port security to be implemented at the application level. NPV makes use of NPIV to get multiple FCIDs allocated from the core switch on the NP port.

The figure below shows a more granular view of an NPV configuration at the interface level.



NPV Mode

A switch is in NPV mode after a user has enabled NPV and the switch has successfully rebooted. NPV mode applies to an entire switch. All end devices connected to a switch that is in NPV mode must log in as an N port to use this feature (loop-attached devices are not supported). All links from the edge switches (in NPV mode) to the NPV core switches are established as NP ports (not E ports), which are used for typical interswitch links. NPIV is used by the switches in NPV mode to log in to multiple end devices that share a link to the NPV core switch.



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 needed and/or configured.

NP Ports

An NP port (proxy N port) is a port on a device that is in NPV mode and connected to the NPV core switch using an F port. NP ports behave like N ports except that in addition to providing N port behavior, they also function as proxies for multiple, physical N ports.

**Note**

A Cisco Nexus 5000 Series switch in NPV mode that runs Cisco NX-OS Release 4.2(1) or later releases supports trunking F port mode on NP ports. You can enable either, or both, VSAN trunking and an F port on an NP port.

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 devices in this NP link are converted to FDISCs. For more details refer to the [Internal FLOGI Parameters, on page 274](#).

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.

Internal FLOGI Parameters

When an NP port comes up, the NPV device first logs itself in to the NPV core switch and sends a FLOGI request that includes the following parameters:

- The fWWN (fabric port WWN) of the NP port used as the pWWN in the internal login.
- The VSAN-based sWWN (switch WWN) of the NPV device used as nWWN (node WWN) in the internal FLOGI.

After completing its FLOGI request, the NPV device registers itself with the fabric name server using the following additional parameters:

- Switch name and interface name (for example, fc1/4) 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 NPV device 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.

The figure below shows the internal FLOGI flows between an NPV core switch and an NPV device.




Table 18: Internal FLOGI Parameters , on page 275 identifies the internal FLOGI parameters that appear in Figure 10-4 .

Table 18: 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 display “switch.” For example, switch: fc1/5.
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:

- Zoning is not enforced at the NPV device (rather, it is enforced on the NPV core switch).
- Multiple devices behind an NPV device log in via the same F port on the core (they use same fWWN and cannot be separated into different zones).
- 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.

Default Port Numbers

Port numbers on NPV-enabled switches will vary depending on the switch model. For details about port numbers for NPV-eligible switches, see the *Cisco NX-OS Family Licensing Guide* .

NPV CFS Distribution over IP

NPV devices use only IP as the transport medium. CFS uses multicast forwarding for CFS distribution. NPV devices do not have ISL connectivity and FC domain. To use CFS over IP, multicast forwarding has to be

enabled on the Ethernet IP switches all along the network that physically connects the NPV switch. You can also manually configure the static IP peers for CFS distribution over IP on NPV-enabled switches. For more information, see the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

NPV Traffic Management

This sections discusses the following aspects of load balancing:

Auto

Before Cisco MDS SAN-OS Release 3.3(1a), NPV supported automatic selection of external links. When a server interface is brought up, an external interface with the minimum load is selected from the available links. There is no manual selection on the server interfaces using the external links. Also, when a new external interface was brought up, the existing load was not distributed automatically to the newly available external interface. This newly brought up interface is used only by the server interfaces that come up after this interface.

Traffic Map

As in Cisco MDS SAN-OS Release 3.3(1a) and NX-OS Release 4.1(1a), NPV supports traffic management by allowing you to select and configure the external interfaces that the server uses to connect to the core switches.



Note

When the NPV traffic management is configured, the server uses only the configured external interfaces. Any other available external interface will not be used.

The NPV traffic management feature provides the following benefits:

- Facilitates traffic engineering by providing dedicated external interfaces for the servers connected to NPV.
- Uses the shortest path by selecting external interfaces per server interface.
- Uses the persistent FC ID feature by providing the same traffic path after a link break, or reboot of the NPV or core switch.
- Balances the load by allowing the user to evenly distribute the load across external interfaces.

Disruptive

Disruptive load balance works independent of automatic selection of interfaces and a configured traffic map of external interfaces. This feature forces reinitialization of the server interfaces to achieve load balance when this feature is enabled and whenever a new external interface comes up. To avoid flapping the server interfaces too often, enable this feature once and then disable it whenever the needed load balance is achieved.

If disruptive load balance is not enabled, you need to manually flap the server interface to move some of the load to a new external interface.

Multiple VSAN Support

By grouping devices into different NPV sessions based on VSANs, it is possible to support multiple VSANs on the NPV-enabled switch. The correct uplink must be selected based on the VSAN that the uplink is carrying.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

NPV Guidelines and Requirements

Following are recommended guidelines and requirements when deploying NPV:

- NPV core switches must support NPIV.
- You can have up to 100 NPV devices.
- Nondisruptive upgrades are supported. See the *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide*.
- Port tracking is supported. See the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.
- 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 or the domain/port of the NPV core switch should be used.
- Port security is supported on the NPV core switch for devices logged in via NPV.
- 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 a specific one to an end device.
- Both servers and targets can be connected to an NPV device.
- Remote SPAN is not supported.
- Local switching is not supported; all traffic is switched using 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.
- In the case of servers that are booted over the SAN with NPV, if an NPV link failover occurs, servers will lose access to their boot LUN temporarily.
- NPV switches do not recognize the BB_SCN configuration on the xNP ports because of interoperability issues with the third-party core switches.

NPV Traffic Management Guidelines

When deploying NPV traffic management, follow these guidelines:

- Use NPV traffic management only when the automatic traffic engineering by the NPV device is not sufficient for the network requirements.
- Do not configure traffic maps for all the servers. For non-configured servers, NPV will use automatic traffic engineering.
- Configure the Persistent FC ID on the core switch. Traffic engineering directs the associated server interface to external interfaces that lead to the same core switch. The server will be assigned the same FC ID for every log in. This guideline is not applicable if a 91x4 switch is used as the core switch.
- Server interfaces configured to a set of external interfaces cannot use any other available external interfaces, even if the configured interfaces are not available.
- Do not configure disruptive load balancing because this involves moving a device from one external interface to another interface. Moving the device between external interfaces requires NPV relogin to the core switch through F port leading to traffic disruption.
- Link a set of servers to a core switch by configuring the server to a set of external interfaces that are linked to the core switch.

DPVM Configuration Guidelines

When NPV is enabled, the following requirements must be met before you configure DPVM on the NPV core switch:

- You must explicitly configure the WWN of the internal FLOGI in DPVM. If DPVM is configured on the NPV core switch for an end device that is connected to the NPV device, then that end device must be configured to be in the same VSAN. Logins from a device connected to an NPV device will fail if the device is configured to be in a different VSAN. To avoid VSAN mismatches, ensure that the internal FLOGI VSAN matches the port VSAN of the NP port.
- The first login from an NP port determines the VSAN of that port. If DPVM is configured for this first login, which is the internal login of the NPV device, then the NPV core switch's VSAN F port is located in that VSAN. Otherwise, the port VSAN remains unchanged.

For details about DPVM configuration, see the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.

NPV and Port Security Configuration Guidelines

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 via NPV, you must adhere to the following requirements:

- The internal FLOGI must be in the port security database so that, the port on the NPV core switch will allow communications and links.
- All of the end device pWWNs must also be in the port security database.

Once these requirements are met, you can enable port security as you would in any other context. For details about enabling port security, see the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.

Configuring N Port Virtualization

Configuring NPV

When you enable NPV, the system configuration is erased and the system reboots with the NPV mode enabled.



Note

We recommend that you save the current configuration either on bootflash or a TFTP server before NPV (if the configuration is required for later use). Use the following commands to save either your non-NPV or NPV configuration: `switch# copy running bootflash:filename` The configuration can be reapplied later using the following command: `switch# copy bootflash:filename running-config`

To use DCNM-SAN and Device Manager to configure NPV, follow these steps:

Procedure

- Step 1** Launch Device Manager from the core NPV switch to enable NPIV on the core NPV switch. From the Admin menu, select **Feature Control**. Select **enable** for the NPIV feature.
- Step 2** Click **Apply**.
- Step 3** From the Interface menu, select **FC All** to configure the NPIV core switch port as an F Port.

- Step 4** In the Mode Admin column, select the **F** port mode and click **Apply**.
- Step 5** Launch Device Manager from the NPV device to enable NPV on the NPV device. From the Admin drop-down menu, select **Feature Control**. Select **enable** for the NPV feature and click **Apply**.
- Step 6** From the Interface drop-down menu, select **FC All** to configure the external interfaces on the NPV device.
- Step 7** In the Mode Admin column, select the **NP** port mode and click **Apply**.
- Step 8** From the Interface drop-down menu, select **FC All** to configure the server interfaces on the NPV device.
- Step 9** In the Mode Admin column, select **F port mode** and click **Apply**.
- Step 10** The default Admin status is **down**. After configuring port modes, you must select up Admin Status to bring up the links.

Using the NPV Setup Wizard

Prerequisites

- For Cisco Nexus 5000 Series switches, you must first enable the NPV mode for the switch by choosing **Switches > N_Port Virtualization (NPV)** in the Physical Attributes pane, and then use the NPV wizard to configure other NPV-related settings on the switch.
- Remove the PortChannel groups if you need to select those particular ports as F ports during the setup. For more information, see the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.

Restrictions

- NPV wizard does not detect ports that are in a channel group and that are not connected by ISLs. The wizard does not configure any port in a PortChannel group to F ports on the core switch. Port channel grouping is not applicable to NPV devices.

Detailed Steps

To configure NPV using the wizard, follow these steps:

Procedure

- Step 1** Select Tools > NPV > NPV Setup... to launch NPV Setup Wizard from DCNM-SAN.

Before the wizard starts, DCNM-SAN checks if there are any NPV- and NPIV-capable switches from the client's SAN. An NPV-capable switch has to be a Cisco MDS 9124, 9134, 9148, a Cisco Nexus 5000 Series switch, an HP Blade Server, or an IBM Blade Server with SAN-OS Release 3.2.2 and later. An NPIV-capable switch has to be Cisco switch with SAN-OS Release 3.0.1 and later. If there are no NPV-capable switches, DCNM-SAN displays an error message saying that no NPV-capable switches are available and that they are not manageable or not present.
- Step 2** Click **OK** to continue.
- Step 3** Select the NPV devices. Click **Next**.

A table lists all the available NPV-capable switches including the switches on which NPV is not yet enabled. Check the check boxes to select the required NPV devices. On devices that are not NPV enabled, this wizard will enable NPV on the devices in the final step.

If you choose switches that are NPV disabled and click Next, a warning message appears with a list of IP addresses of the NPV devices on which NPV will be enabled. Enabling NPV on the switch will result in reboot of the switch. Boot variables of the switches have to be set, to enable NPV on them through this wizard.

Step 4 Select the NPIV core switches. Click **Next**.

Check the check boxes to select the required NPIV core switches. The table lists all the available NPIV core switches including the core switches that have not yet enabled the NPIV feature. NPIV core switches that are not NPIV-enabled. This wizard will enable NPIV in the final step.

Step 5 Create new NPV device and NPIV core switch pairs as required.

Based on selections in the previous steps, the wizard displays all available NPV devices and NPIV core switches in separate lists. You can select one from each list and click Add or Remove buttons to create new NPV device and NPIV core switch combinations or pairs.

The NPV wizard checks if there are any NPIV core switches that are already connected to the NPV devices selected in the previous step. Click the Add Connected Pairs button to add a list of all the existing pairs that are interconnected to the Selected table.

The Selected table is then populated with both the existing and the intended pairs. Each NPIV core switch can be paired with multiple NPV devices.

After Step 6, the wizard prompts you to physically connect the new pairs that are not yet connected.

On the switches that are not paired, the NPV wizard enables the NPV and NPIV modes. However, there is a possibility that these unpaired switches may be segmented and lose their presence on the fabric.

After you click the Next button in Step 3 of 6, the wizard determines if you have selected all the connected pairs. A warning message is displayed that lists all the connected pairs that you have not selected and warns that they will be segmented after the NPV setup.

Step 6 Click **Next**.

Note NPV wizard does not detect ports that are in a channel group and that are not connected by ISLs. The wizard does not configure any port in a Port Channel Group to F ports on the core switch. Port channel grouping is not applicable to NPV devices.

Step 7 You can configure NPV-associated ports either through automated or manual methods.

The Auto Port Selection has two options:

- Choosing the first option allows you to convert the existing ISLs to be run as NPV links. If you want ISLs to take priority, then choose the Convert existing ISLs option.

The wizard discovers ISLs (Up or Down) between the selected switches, that are available at the time of wizard launch.

- Choosing the second option allows the NPV wizard to automatically configure free ports for NPV usage. In the second option, you can choose up to a maximum of six additional NPV links per NPV device and core switch pair.

During automatic port selection on the NPV switch, ports are defined as licensed FC ports with “Operational status” = Auto and “Status Cause” = none(2), offline(8), or sfp not present(29), and “Operational Status” = TE or E.

Ports on the NPV switch are selected in the following way:

The ISLs are considered in the second method. The selection algorithm spreads out the free port selections, so that the first port in every four ports is selected, for example, the 1st, 5th, 9th, etc. If after going through the 1st port in every four ports, you still have not selected enough ports (because the preferred ports were not free) then move to the second port in every four, for example, the 2nd, 6th, 10th etc. Different switches have different port preferences.

Ports on the NPIV switch are selected in the following way:

During automatic port selection on the NPIV switch free ports are defined as ports that are licensed FC ports and ports that have "Operational status" = Auto and "Status Cause" = none(2), offline(8) or sfp not present(29). If the ports are found in any other operational state, (for example F, NP, E, TE etc), then they are considered used, except for E and TE ports that are in ISLs connected to NPV device switches that will be enabled for NPV mode in this wizard session, as they will be considered to be free. However, these ISL ports will not necessarily be the ports selected by the automatic port selection algorithm as they are treated no different than any other free port. If you want to convert those used ISL ports, then choose the Convert existing ISLs option first and then run the wizard a second time choosing Automatic port selection (option 2) to add additional links.

When you choose to configure ports from available ports, the wizard searches for ports that are not currently participating in NP link configuration. It is possible that all ports can be participating in NP port configuration. In that case a warning message is displayed.

Note In both manual and automatic methods of configuring NPV associated ports, the ports that are unhealthy or that are in adminDown state are not considered during port selection.

Select the Manual method to manually create port pairs. Click on a satellite switch and select the NP device port expanded under each of the NPV switches listed. Then select the required F port on the NPIV core switch and click Add for them to pair.

During manual selection from the list for NPV and NPIV, ports are defined as the licensed FC ports with "Operational status" = Auto and "Status Cause" = none(2), offline(8), or sfp not present(29) and "Operational Status" = TE or E.

Note Failed ports with the Auto operational status will not be listed. Failed ports with the E operational status will be listed and available for NPV configuration.

Based on user selection, the wizard decides which ports are set to NP ports on the NPV device side and which are F ports on the core switch side to make an NPV connection.

Note Sometimes the Manual selection in step 4 does not show any port when the NPV switch tree is expanded as the NPV Wizard filters out ports that are in fail or down status. Only healthy ports are made visible in the NPV Switch tree. Check your port settings.

Step 8 Click **Next**.

Step 9 Select a VSAN.

From the drop-down list select a VSAN or enter a VSAN ID to specify the VSAN. All selected NPV devices and NPIV core switches are added to the specified VSAN. All ports on the selected NPV devices and associated ports on the NPIV core switches are added to the VSAN.

Step 10 Click **Next**.

The VSAN configuration is applied in the final step.

Step 11 Review all the NPV Setup configurations you entered in the earlier steps and click **Finish** to complete the setup.

Enable Switch Feature lists the switches, the impending actions against them with reference to features, and the resultant status.

Set Port Type lists the switches and the ports to be set on the switches to configure NPV associate ports.

Configure VSAN lists the switches and ports to be added to the specified VSAN.

Click >> to view the expanded the panes. Click << to collapse the panes.

A progress bar at the bottom of the window indicates the overall extent of completion of the configuration tasks. A text message that runs below the progress bar indicates the current task in progress.

The status cells next to each item indicate the In progress, Success, and Error states. When a configuration cannot be applied, the status cell next to the task is changed to Error. Click Error to view Details. A message is displayed in place of the progress bar stating, Cannot apply all configurations.

After the completion of all the tasks, a View NPV Port Connections link is displayed in the place of the progress bar.

- Step 12** Click View NPV Port Connections to view the NPV port connections in a table. Refer to this list to verify the physical connections between NP Port on NPV devices and Auto ports on NPV core switches. The physical connections already exist for the ISLs and they have to be verified. In some cases when the physical connections do not exist, they have to be established manually.

Configuring NPV Traffic Management


The NPV traffic management feature is enabled after configuring NPV. Configuring NPV traffic management involves configuring a list of external interfaces to the servers, and enabling or disabling disruptive load balancing.

Configuring List of External Interfaces per Server Interface


A list of external interfaces are linked to the server interfaces when the server interface is down, or if the specified external interface list includes the external interface already in use.

To configure the list of external interfaces per server interface, perform the following tasks:

Procedure

- Step 1** Choose **Physical Attributes > Switches > FC Services > N_Port Virtualizer (NPV)**.
- Step 2** Click the **Traffic Map** tab.
- Step 3** Click the  icon in the toolbar or right click and then select **Create Row...**
- Step 4** Select a Switch from the drop-down list.
- Step 5** Type the port numbers or click the [...] button (not available on blade server switches) to select the Server Interface and External Interfaces from the port selection dialog box.

Note You can select only one Server Interface but multiple External Interfaces can be mapped on to it. Previously selected ports are disabled and cannot be selected.

To delete the map entry, select the row from the Traffic Map tab, and then click the  icon in the toolbar or right click and select **Delete Row**.

Enabling the Global Policy for Disruptive Load Balancing

Disruptive load balancing allows you to review the load on all the external interfaces and balance the load disruptively. Disruptive load balancing is done by moving the servers using heavily loaded external interfaces, to the external interfaces running with fewer loads.

To enable or disable the global policy for disruptive load balancing, perform the following tasks:

Procedure

- Step 1** Choose **Physical Attributes > Switches > FC Services > N_Port Virtualizer (NPV)**.
 - Step 2** Click the **Load Balance** tab.
 - Step 3** Check the **Enable** check box to enable disruptive load balancing on the switch.
To enable disruptive load balancing on all the switches, check the **Enable All** check box.
-

Displaying the External Interface Usage for Server Interfaces

To display the external interface usage for the server interfaces, follow these steps:

Procedure

- Step 1** Choose **Physical Attributes > Switches > FC Services > N_Port Virtualizer (NPV)**.
 - Step 2** Click the **External Interface Usage** tab.
-



CHAPTER 11

Configuring Interfaces

- [Configuring Interfaces, on page 285](#)

Configuring Interfaces

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

This chapter includes the following topics:

Information About Interfaces

The main function of a switch is to relay frames from one data link to another. To relay the frames, the characteristics of the interfaces through which the frames are received and sent must be defined. The configured interfaces can be Fibre Channel interfaces, Gigabit Ethernet interfaces, the management interface (mgmt0), or VSAN interfaces.

This section includes the following topics:

Interface Description

For the Fibre Channel interfaces, you can configure the description parameter to provide a recognizable name for the interface. Using a unique name for each interface allows you to quickly identify the interface when you are looking at a listing of multiple interfaces. You can also use the description to identify the traffic or the use for that interface.

Interface Modes

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



Note Interfaces are created in VSAN 1 by default. See the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide* .

Each interface has an associated administrative configuration and an operational status:

- The administrative configuration does not change unless you modify it. This configuration has various attributes that you can configure in administrative mode.
- The operational status represents the current status of a specified attribute like the interface speed. This status cannot be changed and is read-only. Some values may not be valid when the interface is down (for example, the operational speed).

**Note**

When a module is removed and replaced with the same type of module, the configuration is retained. If a different type of module is inserted, then the original configuration is no longer retained.

Each interface is briefly described in the sections that follow.

E Port

In expansion port (E port) mode, an interface functions as a fabric expansion port. This port may be connected to another E port to create an Inter-Switch Link (ISL) between two switches. E ports carry frames between switches for configuration and fabric management. They serve as a conduit between switches for frames destined to remote N ports and NL ports. E ports support class 2, class 3, and class F service.

An E port connected to another switch may also be configured to form a PortChannel (see the *Configuring PortChannels* chapter).

**Note**

We recommend that you configure E ports on 16-port modules. If you must configure an E port on a 32-port oversubscribed module, then you can only use the first port in a group of four ports (for example, ports 1 through 4, 5 through 8, and so forth). The other three ports cannot be used.

F Port

In fabric port (F port) mode, an interface functions as a fabric port. This port may be connected to a peripheral device (host or disk) operating as an N port. An F port can be attached to only one N port. F ports support class 2 and class 3 service.

FL Port

In fabric loop port (FL port) mode, an interface functions as a fabric loop port. This port may be connected to one or more NL ports (including FL ports in other switches) to form a public arbitrated loop. If more than one FL port is detected on the arbitrated loop during initialization, only one FL port becomes operational and the other FL ports enter nonparticipating mode. FL ports support class 2 and class 3 service.

**Note**

FL port mode is not supported on 4-port 10-Gbps switching module interfaces.

NP Ports

An NP port is a port on a device that is in NPV mode and connected to the core switch via an F port. NP ports function like N ports except that in addition to providing N port operations, they also function as proxies for multiple, physical N ports.



Note A Cisco Nexus 5000 Series switch in NPV mode that runs Cisco NX-OS Release 4.2(1) or later releases supports trunking F port mode on NP ports. You can enable either, or both, VSAN trunking and an F port on an NP port.

For more details about NP ports and NPV, see the *Configuring N Port Virtualization* chapter.

TL Port

In translatable loop port (TL port) mode, an interface functions as a translatable loop port. It may be connected to one or more private loop devices (NL ports). TL ports are specific to Cisco MDS 9000 Family switches and have similar properties as FL ports. TL ports enable communication between a private loop device and one of the following devices:

- A device attached to any switch on the fabric
- A device on a public loop anywhere in the fabric
- A device on a different private loop anywhere in the fabric
- A device on the same private loop

TL ports support class 2 and class 3 services.

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop (see the [TL Port ALPA Caches](#), on page 296).



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



Note TL port mode is not supported on Generation 2 switching module interfaces.

TE Port

In trunking E port (TE port) mode, an interface functions as a trunking expansion port. It may be connected to another TE port to create an extended ISL (EISL) between two switches. TE ports are specific to Cisco MDS 9000 Family switches. They expand the functionality of E ports to support the following:

- VSAN trunking
- Transport quality of service (QoS) parameters
- Fibre Channel trace (fctrace) feature

In TE port mode, all frames are transmitted in EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link. This feature is referred to as trunking in the Cisco MDS 9000 Family switches (see the *Configuring Trunking* chapter). TE ports support class 2, class 3, and class F service.

TF Port

In trunking F port (TF port) mode, an interface functions as a trunking expansion port. It may be connected to another trunked N port (TN port) or trunked NP port (TNP port) to create a link between a core switch and

an NPV switch or an HBA to carry tagged frames. TF ports are specific to Cisco MDS 9000 Family switches. They expand the functionality of F ports to support VSAN trunking.

In TF port mode, all frames are transmitted in EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link. This feature is referred to as trunking in the Cisco MDS 9000 Family (see [Chapter 8, “Configuring Trunking.”](#)). TF ports support class 2, class 3, and class F service.

TNP Port

In trunking NP port (TNP port) mode, an interface functions as a trunking expansion port. It may be connected to a trunked F port (TF port) to create a link to a core NPIV switch from an NPV switch to carry tagged frames.

SD Port

In SPAN destination port (SD port) mode, an interface functions as a switched port analyzer (SPAN). The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic that passes 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, they only transmit a copy of the source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports (see the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*).

ST Port

In the SPAN tunnel port (ST port) mode, an interface functions as an entry point port in the source switch for the RSPAN Fibre Channel tunnel. The ST port mode and the remote SPAN (RSPAN) feature are specific to switches in the Cisco MDS 9000 Family. When configured in ST port mode, the interface cannot be attached to any device, and thus cannot be used for normal Fibre Channel traffic (see the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*).



Note

ST port mode is not supported on the Cisco MDS 9124 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

Fx Port

Interfaces configured as Fx ports can operate in either F port or FL port mode. The Fx port mode is determined during interface initialization depending on the attached N port or NL port. This administrative configuration disallows interfaces to operate in any other mode—for example, preventing an interface to connect to another switch.

B Port

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as the Cisco PA-FC-1G Fibre Channel port adapter, implement a bridge port (B port) model to connect geographically dispersed fabrics. This model uses B ports as described in the T11 Standard FC-BB-2.

If an FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled (see the *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*).

Auto Mode

Interfaces configured in auto mode can operate in one of the following modes: F port, FL port, E port, TE port, or TF port. The port mode is determined during interface initialization. For example, if the interface is connected to a node (host or disk), it operates in F port or FL port mode depending on the N port or NL port mode. If the interface is attached to a third-party switch, it operates in E port mode. If the interface is attached to another switch in the Cisco MDS 9000 Family, it may become operational in TE port mode (see the *Configuring Trunking* chapter).

TL ports and SD ports are not determined during initialization and are administratively configured.



Note

Fibre Channel interfaces on Storage Services Modules (SSMs) cannot be configured in auto mode.

Interface States

The interface state depends on the administrative configuration of the interface and the dynamic state of the physical link.

Administrative States

The administrative state refers to the administrative configuration of the interface as described in below table.

Table 19: Administrative States

Administrative State	Description
Up	Interface is enabled.
Down	Interface is disabled. If you administratively disable an interface by shutting down that interface, the physical link layer state change is ignored.

Operational States

The operational state indicates the current operational state of the interface as described in below table.

Table 20: 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 or TF mode.

Reason Codes

Reason codes are dependent on the operational state of the interface as described in below table.

Table 21: 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 22: Reason Codes for Nonoperational States , on page 291.



Note Only some of the reason codes are listed in [Table 22: Reason Codes for Nonoperational States](#) , on page 291.

If the administrative state is up and the operational state is down, the reason code differs based on the nonoperational reason code as described in below table.

Table 22: 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.	
Reconfigure fabric in progress	The fabric is currently being reconfigured.	
Offline	The Cisco NX-OS software waits for the specified R_A_TOV time before retrying initialization.	
Inactive	The interface VSAN is deleted or is in a suspended state. To make the interface operational, assign that port to a configured and active VSAN.	
Hardware failure	A hardware failure is detected.	
Error disabled	Error conditions require administrative attention. Interfaces may be error-disabled for various reasons. For example: <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 next, administratively shut down or enable the interface.	
FC redirect failure	A port is isolated because a Fibre Channel redirect is unable to program routes.	
No port activation license available	A port is not active because it does not have a port license.	
SDM failure	A port is isolated because SDM is unable to program routes.	

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.	
Nonparticipating	FL ports cannot participate in loop operations. It may happen if more than one FL port exists in the same loop, in which case all but one FL port in that loop automatically enters nonparticipating mode.	Only FL ports and TL ports
PortChannel administratively down	The interfaces belonging to the PortChannel are down.	Only PortChannel interfaces
Suspended due to incompatible speed	The interfaces belonging to the PortChannel have incompatible speeds.	
Suspended due to incompatible mode	The interfaces belonging to the PortChannel have incompatible modes.	
Suspended due to incompatible remote switch WWN	An improper connection is detected. All interfaces in a PortChannel must be connected to the same pair of switches.	

Graceful Shutdown

Interfaces on a port are shut down by default (unless you modified the initial configuration).

The Cisco NX-OS software implicitly performs a graceful shutdown in response to either of the following actions for interfaces operating in the E port mode:

- If you shut down an interface.
- If a Cisco NX-OS software application executes a port shutdown as part of its function.

A graceful shutdown ensures that no frames are lost when the interface is shutting down. When a shutdown is triggered either by you or the Cisco NX-OS software, the switches connected to the shutdown link coordinate

with each other to ensure that all frames in the ports are safely sent through the link before shutting down. This enhancement reduces the chance of frame loss.

A graceful shutdown is not possible in the following situations:

- If you physically remove the port from the switch.
- If in-order delivery (IOD) is enabled (for information about IOD, refer to the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*).
- If the Min_LS_interval interval is higher than 10 seconds. For information about FSPF global configuration, refer to the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.



Note This feature is only triggered if both switches at either end of this E port interface are MDS switches and are running Cisco SAN-OS Release 2.0(1b) or later, or MDS NX-OS Release 4.1(1a) or later.

Port Administrative Speeds

By default, the port administrative speed for an interface is automatically calculated by the switch.

For internal ports on the Cisco Fabric Switch for HP c-Class BladeSystem and Cisco Fabric Switch for IBM BladeCenter, a port speed of 1 Gbps is not supported. Auto-negotiation is supported between 2 Gbps and 4 Gbps only. Also, if the BladeCenter is a T chassis, then port speeds are fixed at 2 Gbps and auto-negotiation is not enabled.

Autosensing

Autosensing speed is enabled on all 4-Gbps and 8-Gbps switching module interfaces by default. This configuration enables the interfaces to operate at speeds of 1 Gbps, 2 Gbps, or 4 Gbps on the 4-Gbps switching modules, and 8 Gbps on the 8-Gbps switching modules. 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.

To avoid wasting unused bandwidth on 48-port and 24-port 4-Gbps and 8-Gbps Fibre Channel switching modules, you can specify that only 2 Gbps of required bandwidth be reserved, not the default of 4 Gbps or 8 Gbps. This feature shares the unused bandwidth within the port group provided that it does not exceed the rate limit configuration for the port. You can also use this feature for shared rate ports that are configured for autosensing.



Tip When migrating a host that supports up to 2-Gbps traffic (that is, not 4 Gbps with autosensing capabilities) to the 4-Gbps switching modules, use autosensing with a maximum bandwidth of 2 Gbps. When migrating a host that supports up to 4-Gbps traffic (that is, not 8 Gbps with autosensing capabilities) to the 8-Gbps switching modules, use autosensing with a maximum bandwidth of 4 Gbps.

Frame Encapsulation

The **switchport encap eisl** command only applies to SD port interfaces. This command determines the frame format for all frames transmitted by the interface in SD port mode. If the encapsulation is set to EISL, all outgoing frames are transmitted in the EISL frame format, regardless of the SPAN sources.

The **switchport encap eisl** command is disabled by default. If you enable encapsulation, all outgoing frames are encapsulated, and you will see a new line (Encapsulation is eisl) in the **show interface SD_port_interface** command output. See the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide* .

You can set the frame format to EISL for all frames transmitted by the interface in SD port mode. If you sent the frame encapsulation to EISL, all outgoing frames are transmitted in the EISL frame format, regardless of the SPAN sources. See the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide* .

Refer to the Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide to configure frame encapsulation on an interface.

Beacon LEDs

The following figure displays the status, link, and speed LEDs in a 16-port switching module.

1	Status LED ⁴	3	Link LEDs 1 and speed LEDs ⁵
2	1/2-Gbps Fibre Channel port group ⁶	4	Asset tag ⁷

⁴ See the Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide .

⁵ See the “Speed LEDs” section.

⁶ See the Generation 1 Interface Configuration Guidelines section.

⁷ Refer to the Cisco MDS 9000 Family hardware installation guide for your platform .

Speed LEDs

Each port has one link LED on the left and one speed LED on the right.

The speed LED displays the speed of the port interface:

- Off—The interface attached to that port is functioning at 1000 Mbps.
- On (solid green)—The interface attached to that port is functioning at 2000 Mbps (for 2 Gbps interfaces).

The speed LED also displays if the beacon mode is enabled or disabled:

- Off or solid green—Beacon mode is disabled.
- Flashing green—The beacon mode is enabled. The LED flashes at one-second intervals.



Note

Generation 2, Generation 3, and Generation 4 modules and fabric switches do not have speed LEDs.

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.
- Momentary sync loss.

- Loose cable connection at one or both ends.
- Improper GBIC or SFP connection at one or both ends.

A bit error rate threshold is detected when 15 error bursts occur in a 5-minute period. By default, the switch disables the interface when the threshold is reached. You can enter a **shutdown** and **no shutdown** command sequence to re-enable the interface.

You can configure the switch to not disable an interface when the threshold is crossed. By default, the threshold disables the interface.

SFP Transmitter Types

The small form-factor pluggable (SFP) hardware transmitters are identified by their acronyms when displayed. [Table 23: SFP Transmitter Acronym Definitions](#), on page 295 defines the acronyms used for SFPs.

The small form-factor pluggable (SFP) hardware transmitters are identified by their acronyms when displayed in the **show interface brief** command. If the related SFP has a Cisco-assigned extended ID, then the **show interface** and **show interface brief** commands display the ID instead of the transmitter type. The **show interface transceiver** command and the **show interface fc slot/port transceiver** command display both values for Cisco-supported SFPs. Below table defines the acronyms used in the command output.

Table 23: SFP Transmitter Acronym Definitions

Definition	Acronym
Standard transmitters defined in the GBIC specifications	
short wave laser	swl
long wave laser	lwl
long wave laser cost reduced	lwcr
electrical	elec
Extended transmitters assigned to Cisco-supported SFPs	
CWDM-1470	c1470
CWDM-1490	c1490
CWDM-1510	c1510
CWDM-1530	c1530
CWDM-1550	c1550
CWDM-1570	c1570
CWDM-1590	c1590
CWDM-1610	c1610

See the [Displaying SFP Transmitter Types](#), on page 309.

TL Ports

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop. The legacy devices are used in Fibre Channel networks, and devices outside the loop may need to communicate with them. The communication functionality is provided through TL ports. See the [Interface Modes, on page 285](#).

TL port mode is not supported on the following hardware:

- Generation 2 switching module interfaces
- Cisco MDS 9124 Fabric Switch
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter

Below table lists the TL port translations supported in Cisco MDS 9000 Family switches.

Table 24: Supported TL Port Translations

Translation from	Translation to	Example
Private initiator	Private target	From I1 to T1 or vice versa
Private initiator	Public target — N port	From I1 to T2 or vice versa
Private initiator	Public target — NL port	From I4 to T3 or vice versa
Public initiator — N port	Private target	From I2 to T1 or vice versa
Public initiator — NL port	Private target	From I3 to T1 or vice versa

The followin figure shows examples of TL port translation support.

TL Port ALPA Caches

Although TL ports cannot be automatically configured, you can manually configure entries in arbitrated loop physical address (ALPA) caches. Generally, ALPA cache entries are automatically populated when an ALPA is assigned to a device. Each device is identified by its port world wide name (pWWN). When a device is allocated an ALPA, an entry for that device is automatically created in the ALPA cache.

A cache contains entries for recently allocated ALPA values. These caches are maintained on various TL ports. If a device already has an ALPA, the Cisco NX-OS software attempts to allocate the same ALPA to the device each time. The ALPA cache is maintained in persistent storage and saves information across switch reboots. The maximum cache size is 1000 entries. If the cache is full, and a new ALPA is allocated, the Cisco NX-OS software discards an inactive cache entry (if available) to make space for the new entry. See the [TL Port, on page 287](#) for more information on TL ports.

Refer to the Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide to manage the TL Port ALPA cache.

Port Guard

The port guard feature is intended for use in environments where the system and application environment does not adapt quickly and efficiently to a port going down and back up, or to a port rapidly cycling up and down, which can happen in some failure modes. For example, if a system takes five seconds to stabilize after

a port goes down, but the port is going up and down once a second, a more severe failure in the fabric might occur.

The port guard feature gives the SAN administrator the ability to prevent this issue from occurring in environments that are vulnerable to these problems. The port can be configured to stay down after the first failure or after a specified number of failures in a specified time period. This allows the SAN administrator to intervene and control the recovery, avoiding any problems caused by the cycling.

Using the port guard feature, you can restrict the number of error reports and bring a malfunctioning port to down state dynamically. A port can be configured to go into error-disabled state for specific types of failures.

A general link failure caused by link-down is the superset of all other causes. The sum of the number of all other causes equals to the number of link-down link failures. This means a port is brought to down state when it reaches the maximum number of allowed link failures or the number of specific causes.

The causes of link failure can be any of the following:

- ESP trustsec-violation
- Bit-errors
- Signal loss
- Sync loss
- Link reset
- Credit loss
- Additional causes might be the following:
 - Not operational (NOS).
 - Too many interrupts.
 - Cable is disconnected.
 - Hardware recoverable errors.
 - The connected device rebooted (F ports only).
 - The connected linecard rebooted (ISL only).

Port Monitor

Port monitor helps to monitor the performance and the status of ports and generate alerts when problems occur. You can configure the thresholds for various counters and trigger an event when the values cross the threshold settings.

The default port monitor policy has the following threshold values:

Counter	Threshold Type	Interval (Seconds)	% Rising Threshold	Event	% Falling Threshold	Event
Link Loss	Delta	60	5	4	1	4
Sync Loss	Delta	60	5	4	1	4
Protocol Error	Delta	60	1	4	0	4
Signal Loss	Delta	60	5	4	1	4
Invalid Words	Delta	60	1	4	0	4
Invalid CRCs	Delta	60	5	4	1	4

Counter	Threshold Type	Interval (Seconds)	% Rising Threshold	Event	% Falling Threshold	Event
RX Performance	Delta	60	2147483648	4	524288000	4
TX Performance	Delta	60	2147483648	4	524288000	4

Port Monitor Port Guard

Port monitor port guard is a feature that disables or shuts down a port when an event occurs. Depending on the configuration, when an event occurs the port is either error-disabled or flapped.

Port monitor port guard is a different or separate feature that functions based on the configuration of the **errordisable** command.

Port Group Monitor

Each line card or module has a predefined set of ports which share the same backplane bandwidth called port groups. While oversubscription is a feature, the port group monitor feature helps to monitor the spine bandwidth utilization. An alarm syslog is generated so that you can provision the ports across port groups evenly to manage the oversubscription better.

When the port group monitor feature is enabled and a policy consisting of polling interval in seconds, and the raising and falling thresholds in percentage are specified, port group monitor generates a syslog if a port group traffic goes above the specified percentage of the maximum supported bandwidth for that port group (for rx and for tx) and another syslog if the value falls below the specified threshold.

The default port group policy has the following threshold values:

Counter	Threshold Type	Interval (Seconds)	% Rising Threshold	% Falling Threshold
RX Performance	Delta	60	80	20
TX Performance	Delta	60	80	20

Local Switching

Local switching can be enabled in Generation 4 modules, which allows traffic to be switched directly with a local crossbar when the traffic is directed from one port to another on the same line card. By using local switching, an extra switching step is avoided, which decreases the latency.

When using local switching, note the following guidelines:

- All ports need to be in shared mode, which usually is the default state. To place a port in shared mode, enter the **switchport ratemode shared** command.
- E ports are not allowed in the module because they must be in dedicated mode.

Slow Drain Device Detection and Congestion Avoidance

All data traffic between end devices in a SAN fabric is carried by Fibre Channel Class 3. In some cases, the traffic is carried by Class 2 services that use link-level, per-hop-based, and buffer-to-buffer flow control.

These classes of service do not support end-to-end flow control. When there are slow devices attached to the fabric, the end devices do not accept the frames at the configured or negotiated rate. The slow devices lead to ISL credit shortage in the traffic destined for these devices and they congest the links. The credit shortage affects the unrelated flows in the fabric that use the same ISL link even though destination devices do not experience slow drain.

This feature provides various enhancements to detect slow drain devices that are causing congestion in the network and also provides a congestion avoidance function.

This feature is focused mainly on the edge ports that are connected to slow drain devices. The goal is to avoid or minimize the frames being stuck in the edge ports due to slow drain devices that are causing ISL blockage. To avoid or minimize the stuck condition, configure lesser frame timeout for the ports. No-credit timeout drops all packets once the slow drain is detected using the configured thresholds. The lesser frame timeout value helps to alleviate the slow drain condition that affects the fabric by dropping the packets on the edge ports sooner than the time they actually get timed out (500 ms). This function frees the buffer space in ISL, which can be used by other unrelated flows that do not experience slow drain condition.

**Note**

This feature is used mainly for edge ports that are connected to slow edge devices. Even though this feature can be applied to ISLs as well, we recommend that you apply this feature only for edge F ports and retain the default configuration for ISLs as E and TE ports. This feature is not supported on Generation 1 modules.

Management Interfaces

You can remotely configure the switch through the management interface (mgmt0). To configure a connection on the mgmt0 interface, you must configure either the IP version 4 (IPv4) parameters (IP address, subnet mask, and default gateway) or the IP version 6 (IPv6) parameters so that the switch is reachable.

Before you begin to configure the management interface manually, obtain the switch's IPv4 address and subnet mask, or the IPv6 address.

The management port (mgmt0) is autosensing and operates in full-duplex mode at a speed of 10/100/1000 Mbps. Autosensing supports both the speed and the duplex mode. On a Supervisor-1 module, the default speed is 100 Mbps and the default duplex mode is auto. On a Supervisor-2 module, the default speed is auto and the default duplex mode is auto.

**Note**

You need to explicitly configure a default gateway to connect to the switch and send IP packets or add a route for each subnet.

VSAN Interfaces

VSANs apply to Fibre Channel fabrics and enable you to configure multiple isolated SAN topologies within the same physical infrastructure. You can create an IP interface on top of a VSAN and then use this interface to send frames to this VSAN. To use this feature, you must configure the IP address for this VSAN. VSAN interfaces cannot be created for nonexistent VSANs.

Prerequisites for Interfaces

Before you begin configuring the interfaces, ensure that the modules in the chassis are functioning as designed. To verify the status of a module at any time, enter the **show module** command in EXEC mode. For information about verifying the module status, refer to the *Cisco NX-OS Fundamentals Configuration Guide*.

Guidelines and Limitations

This section includes the following topics:

Generation 1 Interface Configuration Guidelines

The Generation 1 interfaces configuration guidelines apply to the following hardware:

- The 32-port, 2-Gbps or 1-Gbps switching module interfaces
- The Cisco MDS 9140 and 9120 switch interfaces



Note

Due to the hardware design of the MDS 9134 switch, we do not support interface out-of-service action on either of its two 10-Gigabit ports. This is because no internal port hardware resource is released when an out-of-service action is performed on these 10-Gigabit ports.

When configuring these host-optimized ports, the following port mode guidelines apply:

- You can configure only the first port in each 4-port group (for example, the first port in ports 1-4, the fifth port in ports 5-8, and so on) as an E port. If the first port in the group is configured as an E port, the other three ports in each group (ports 2-4, 6-8, and so on) are not usable and remain shutdown.
- If you execute the **write erase** command on a 32-port switching module, and then copy a saved configuration to the switch from a text file that contains the **no system default switchport shutdown** command, you need to copy the text file to the switch again for the E ports to come up without manual configuration.
- If any of the other three ports are enabled, you cannot configure the first port as an E port. The other three ports continue to remain enabled.
- The auto mode is not allowed in a 32-port switching module or the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).
- The default port mode is Fx (Fx negotiates to F or FL) for 32-port switching modules.
- The 32-port switching module does not support FICON.



Note

We recommend that you configure your E ports on a 16-port switching module. If you must configure an E port on a 32-port host-optimized switching module, the other three ports in that 4-port group cannot be used.



Note

In the Cisco MDS 9100 Series, the groups of ports that are located on the left and outlined in white are full line rate. The other ports are host-optimized. Each group of 4 host-optimized ports have the same features as for the 32-port switching module.

Private Loop Configuration Guidelines

Follow these guidelines when configuring private loops:

- A maximum of 64 fabric devices can be proxy to a private loop.
- Fabric devices must be in the same zone as private loop devices to be proxy to the private loop.
- Each private device on a TL port may be included in a different zone.
- All devices on the loop are treated as private loops. You cannot mix private and public devices on the loop if the configured port mode is TL.
- The only FC4-type supported by TL ports is SCSI (FCP).
- Communication between a private initiator to a private target on the same private loop does not invoke TL port services.

VSAN Interface Configuration Guidelines

Follow these guidelines when creating or deleting VSAN interfaces:

- Create a VSAN before creating the interface for that VSAN. If a VSAN does not exist, the interface cannot be created.
- Create the interface VSAN—it is not created automatically.
- If you delete the VSAN, the attached interface is automatically deleted.
- Configure each interface only in one VSAN.



Tip After configuring the VSAN interface, you can configure an IP address or Virtual Router Redundancy Protocol (VRRP) feature. See the *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide* .

Default Settings

[Table 25: Default Interface Parameters](#) , on page 301 lists the default settings for interface parameters.

Table 25: Default Interface Parameters

Parameters	Default
Interface mode	Auto
Interface speed	Auto
Administrative state	Shutdown (unless changed during initial setup)
Trunk mode	On (unless changed during initial setup) on non-NPV and NPIV core switches. Off on NPV switches.
Trunk-allowed VSANs or VF-IDs	1 to 4093
Interface VSAN	Default VSAN (1)
Beacon mode	Off (disabled)
EISL encapsulation	Disabled

Parameters	Default
Data field size	2112 bytes

Configuring Interfaces

This section includes the following topics:

For more information on configuring mgmt0 interfaces, refer to the *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide* and *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*.

For more information on configuring Gigabit Ethernet interfaces, see the *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*.

Common Interface Configuration

Some configuration settings are similar for Fibre Channel, management, and VSAN interfaces. You can configure interfaces from DCNM-SAN by expanding Switches > FC Interfaces and selecting either the Physical or Logical interface type from the Physical Attributes pane.

Setting the Interface Administrative State

To disable or enable an interface using DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Expand Switches > FC Interfaces > Physical. You see the interface configuration in the Information pane.
 - Step 2** Click the General tab.
 - Step 3** Click Mode admin.
You see the drop-down box.
 - Step 4** Set the status or mode to the required status.
 - Step 5** (Optional) Set other configuration parameters using the other tabs.
 - Step 6** Click Apply Changes.
-

Configuring Interface Modes

To configure the interface mode using DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Expand Switches > FC Interfaces > Physical.
You see the interface configuration in the Information pane.
 - Step 2** Click the General tab.
 - Step 3** Click Mode Admin. Set the desired interface mode from the Admin drop-down menu.

- Step 4** (Optional) Set other configuration parameters using the other tabs.
- Step 5** Click Apply Changes icon.

Configuring 10-Gbps FC Mode

The 48-port 8-Gbps Advanced Fibre Channel module (DS-X9248-256K9) and the 32-port 8-Gbps Advanced Fibre Channel module (DS-X9232-256K9) can switch between two speed modes—the 1-, 2-, 4-, 8-Gbps or 10-Gbps. By default, the modules are online in the 1-, 2-, 4-, and 8-Gbps modes when they are loaded for the first time. There are two ways to change the ports to the 10-Gbps speed mode:

- Using the **10G-speed mode** command, which is the recommended method.
- Using the generic speed configuration **switchport speed** command which has certain constraints.

The following conditions apply when the ports in the module can be configured to 10-Gbps speed mode:

- The ports in the module can be configured to 10-Gbps speed only when the DS-13SLT-FAB3 module bandwidth is 256-G. Any other combination of fabric modules will not let the ports come up in 10-Gbps.
- When in 10-Gbps mode, the ports in the module that are not 10-Gbps capable are disabled and will be in out-of-service state.
- The ports function only in full rate mode. They cannot be moved to shared rate mode.
- The ports cannot be configured in any other speed.
- Ports that are capable of 10-Gbps that are disabled or out-of-service cannot be put back in service using the **no out-of-service** command. To put these ports back in service, all ports in the module first have to be moved to the out-of-service state. Then they can be brought back to the in service state.
- Local switching must be disabled, otherwise, ports cannot be configured in dedicated mode.

Only certain ports on the 48-port and 32-port 8-Gbps Advanced Fibre Channel modules are 10-Gbps capable. When running in 10-Gbps mode, the non-10-Gbps ports cannot be operational. They have to be either in shut state or out-of-service state.

To configure the interface mode using DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Expand Switches > FC Interfaces > Physical.
- You see the interface configuration in the Information pane.
- Step 2** Click the General tab.
- Step 3** Click Mode Admin. Set the desired interface mode from the Admin drop-down menu.
- Step 4** (Optional) Set other configuration parameters using the other tabs.
- Step 5** Click Apply Changes icon.
-

Configuring Port Administrative Speeds



Note Changing the port administrative speed is a disruptive operation.

To configure the administrative speed of the interface using DCNM-SAN, follow these steps:

Procedure

Step 1 Expand Switches > FC Interfaces > Physical.

You see the interface configuration in the Information pane.

Step 2 Click the General tab.

Step 3 Click Speed Admin. Set the desired speed from the drop-down menu.

The number indicates the speed in megabits per second (Mbps). You can set the speed to 1-Gbps, 2-Gbps, 4-Gbps, 8-Gbps, autoMax2G, autoMax4G, or auto (default).

Note On a Cisco Nexus 5000 Series switch that runs Cisco NX-OS Release 4.2(2), you can configure the 8-Gbps administrative speed only on a M1060 switch module. You can configure the speed to 1-Gbps, 2-Gbps, or 4-Gbps on all switch modules on a Cisco Nexus 5000 Series switch that runs Cisco NX-OS Release 4.2(2) or earlier releases.

Step 4 Click Apply Changes.

Configuring Port Speed Group

To configure the administrative speed of the interface using DCNM-SAN, follow these steps:

Procedure

Step 1 Expand Switches > FC Interfaces > Physical.

You see the interface configuration in the Information pane.

Step 2 Click the General tab.

Step 3 Click SpeedGroup. Set the desired speed group from the drop-down menu.

You can select any of the speed groups from the menu list—notApplicable, tenG, oneTwoFourEightG, or twoFourEightSixteenG.

Note For a DS-X9248-256K9 or DS-X9232-256K9 line card, the speed group must be set to tenG.

Step 4 Click Apply Changes.

Specifying a Port Owner Using DCNM-SAN

To specify or remove the port owner using DCNM-SAN, follow these steps:

Procedure

Step 1 Expand Switches > FC Interfaces > Physical.

You see the interface configuration in the Information pane.

- Step 2** Click the **General** tab and then select the switch/port.
 - Step 3** In the Owner text box, enter a port owner and the purpose for which port is used.
-

Specifying a Port Owner Using Device Manager

To specify or remove the port owner using Device Manager, follow these steps:

Procedure

- Step 1** Double-click the interface in the modules panel.
 - Step 2** Click the **General** tab.
 - Step 3** In the Owner text box, enter a port owner and the purpose for which the port is used.
 - Step 4** Click **Apply**.
-

Configuring Beacon Mode

By default, the beacon mode is disabled on all switches. The beacon mode is indicated by a flashing green light that helps you identify the physical location of the specified interface. Configuring the beacon mode has no effect on the operation of the interface.

To enable beacon mode for a specified interface or range of interfaces using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand **Switches > Ethernet Interfaces > Physical > IPS (the Gigabit Ethernet Interfaces)**.
You see the interface configuration in the Information pane.
 - Step 2** Click the Beacon Mode and enable this option for the selected switch.
 - Step 3** Click **Apply Changes**.
-

Troubleshooting Tips

The flashing green light turns on automatically when an external loopback is detected that causes the interfaces to be isolated. The flashing green light overrides the beacon mode configuration. The state of the LED is restored to reflect the beacon mode configuration after the external loopback is removed.

Configuring TL Ports

Private loops require setting the interface mode to TL.

Use the **switchport mode** command to configure a TL port. See the [Configuring Interface Modes, on page 302](#).

To configure the TL interface mode using DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Expand Switches > FC Interfaces > Physical. You see the interface configuration in the Information pane.
 - Step 2** Click the General tab and click Mode Admin.
 - Step 3** Set the Mode Admin drop-down menu to the required status.
 - Step 4** (Optional) Set other configuration parameters using the other tabs.
 - Step 5** Click Apply Changes.
-

Configuring Port Guard Using DCNM-SAN

To enable port guard using DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Expand Switches > FC Interfaces > Physical > Port Guard from the Physical Attributes pane.
You see the interfaces listed in the Information pane.
 - Step 2** Click the Link Down tab and then select a switch or port.
 - Step 3** Check the check box in the Enable column.
 - Step 4** (Optional) Enter the Duration in seconds and the number of flaps. If the values are 0, the port is brought to down state if the link flaps even once. Otherwise, the link is brought to down state if the link flaps for the number of flaps within the duration.
 - Step 5** Click Apply Changes to activate the configuration.
 - Step 6** Click the **TrustSec Violation** tab, and then select a switch or port.
 - Step 7** Check the check box in the Enable column.
 - Step 8** (Optional) Enter the duration in seconds and the number of flaps. If the values are 0, the port is brought to down state if a trustsec violation occurs even once. Otherwise, the link is brought to down state if there is trustsec violation for the number of flaps within the duration.
 - Step 9** Click the **Bit Errors**, **Signal Loss**, **Sync Loss**, **Link-reset**, and **Credit Loss** tabs and complete the port guard configuration.
 - Step 10** Click Apply Changes to activate the configuration.
-

Configuring Port Guard Using Device Manager

To enable port guard for single or multiple interfaces using Device Manager, follow these steps:

Procedure

-
- Step 1** Expand Switches > FC Interfaces > Physical > Port Guard from the Physical Attributes pane.

You see the FC Interfaces listed.

- Step 2** Click the Link Down tab, and then select the switch or port.
- Step 3** Check the check box in the Enable column.
- Step 4** (Optional) Enter the duration in seconds and the number of flaps. If the values are 0, the port goes into a down state even if the link flaps once. Otherwise, the link goes into a down state if the link flaps for the number of flaps within the duration.
- Step 5** Click Apply Changes to activate the configuration.
- Step 6** Click the **TrustSec Violation** tab, and then select the switch or port.
- Step 7** Check the check box in the Enable column.
- Step 8** (Optional) Enter the Duration in seconds and the number of flaps. If the values are 0, the port is brought to down state if a trustsec violation occurs even once. Otherwise, the link is brought to down state if a trustsec violation occurs for the number of flaps within the duration.
- Step 9** Click Apply Changes to activate the configuration.

Troubleshooting Tips

- Link down is the superset of all other causes. A port is brought to down state if the total number of other causes equals to the number of allowed link-down failures.
- Even if the link does not flap due to failure of the link, and port guard is not enabled, the port goes into a down state if too many invalid FLOGI requests are received from the same host. Use the **shut** and the **no shut** commands consecutively to bring up the link.

Configuring Management Interfaces

To configure the management interface using DCNM-SAN, follow these steps:

Procedure

- Step 1** Select a VSAN in the Logical Domains pane.
- Step 2** In the Physical Attributes pane, expand Switches > Management.
You see the interface configuration in the Information pane.
- Step 3** Click the **IP Addresses** tab and set the Interface, IP Address/Mask field.
- Step 4** Click the General tab and set the Status Admin to up.
- Step 5** (Optional) Set other configuration parameters using the other tabs.
- Step 6** Click Apply Changes.

Creating VSAN Interfaces

To create a VSAN interface using DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Expand Switches > Management.
- Step 2** Click Create Row.
You see the Create Interface dialog box.
- Step 3** Select the switch and VSAN ID for which you want to configure a VSAN interface.
- Note** You can only create a VSAN interface for an existing VSAN. If the VSAN does not exist, you cannot create a VSAN interface for it.
- Step 4** Set IPAddress/Mask to the IP address and subnet mask for the new VSAN interface.
- Step 5** Click Create to create the VSAN interface or click Close to close the dialog box without creating the VSAN interface.
-

Configuring Average Credit Nonavailable Duration Threshold and Action

When the average credit nonavailable duration exceeds the set threshold, the port can be error-disabled, a trap can be sent with interface details, and a syslog can be generated with interface details. One or more of these actions can also be combined together. These actions can be turned on or off depending on the situation. The port monitor feature provides the command line interface to configure the thresholds and action. The threshold configuration can be a percentage of credit nonavailable duration in an interval.

The thresholds are that the credit nonavailable duration can be 0 percent to 100 percent in multiples of 10, and the interval can be 1 second to 1 hour. The default is 10 percent in 1 second and generates a syslog.

To configure average credit nonavailable duration threshold and action, refer to the [Port Monitor, on page 297](#).



Note This feature is not supported on 1 RU fabric switches.

Verifying Interfaces Configuration

This section includes the following topics:

Displaying the Owned Ports

To display the interfaces owned using Device Manager, follow these steps:

Procedure

-
- Step 1** From the menu bar, click the **Ports All** drop-down button.
- Step 2** From the drop-down list, select **Owned**.
-

Obtaining Interface Statistics

You can use DCNM-SAN or Device Manager to collect interface statistics on any switch. These statistics are collected at intervals that you can set.



Note In DCNM-SAN, you can collect interface statistics by expanding ISLs and selecting Statistics from the Physical Attributes pane.

To obtain and display interface counters using Device Manager, follow these steps:

Procedure

- Step 1** From the menu bar, click **Interface**. Select Monitor.
Select any of the Interfaces that are displayed. For example, Virtual FC Enabled.
- 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** Select any tab to view those related statistics.
- 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 select the Print icon to print the statistics.
- Step 6** Click Close when you are finished gathering and displaying statistics.

Displaying SFP Transmitter Types

To show the SFP types for an interface using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand Switches > FC Interfaces > 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.

Monitoring a Port Group

To monitor a particular group using Device Manager, follow these steps:

Procedure

- Step 1** Right-click any port group module and select **Check Oversubscription**.
The **Check Oversubscription** table is displayed.

Step 2 From the **Monitor** drop-down list, select one particular group to monitor.

The Device Manager displays the monitoring table of the selected group with counters on each interval and displays the line chart automatically. From the Monitoring table, you can also choose the **Bar chart** icon to view the selected group as bar charts.



CHAPTER 12

Configuration of Fibre Channel Interfaces

- [Configuration of Fibre Channel Interfaces, on page 311](#)

Configuration of Fibre Channel Interfaces

This chapter describes how to configure the Fibre Channel interfaces.

This chapter includes the following topics:

Information About Fibre Channel Interfaces

This section includes the following topics:

Generations of Modules and Switches

Cisco MDS 9000 Family hardware modules and switches are categorized into generations based on the time of introduction, capabilities, features, and compatibilities:

- Generation 1—Modules and switches with a maximum port speed of 2 Gbps.
- Generation 2—Modules and switches with a maximum port speed of 4 Gbps.
- Generation 3—Modules and switches with a maximum port speed of 8 Gbps.
- Generation 4—Modules with a maximum port speed of 8-Gbps or 10-Gbps.

The Cisco MDS 9500 Series switches, Cisco MDS 9222i, Cisco MDS 9216A, and Cisco MDS 9216i switches support the Generation 2 modules. Each module or switch can have one or more ports in port groups that share common resources such as bandwidth and buffer credits.

In addition to supporting Generation 2 modules, the Cisco MDS 9500 Series switches and the Cisco MDS 9222i switch support the Generation 3 modules. Similar to Generation 2, each Generation 3 or Generation 4 module can have one or more ports in port groups that share common resources such as bandwidth and buffer credits.

Generation 3 modules are supported on the Cisco MDS 9506 and 9509 switches with Supervisor-2 modules. The MDS 9513 Director supports 4/44-port Host-Optimized Fibre Channel switching module with either Fabric 1 or Fabric 2 modules, but requires Fabric 2 module for support of the 48-port and the 24-port 8-Gbps Fibre Channel switching modules. The MDS 9222i switch supports the 4/44-port Host-Optimized Fibre Channel switching module.

The Cisco 9500 Series switches support the following Generation 4 modules: the 48-port 8-Gbps Advanced Fibre Channel switching module (DS-X9248-256K9) and the 32-port 8-Gbps Advanced Fibre Channel module (DS-X9232-256K9). Cisco MDS NX-OS Release 6.x or higher is required to support the Generation 4 modules.

[Table 26: Fibre Channel Modules and Fabric Switches, on page 312](#) identifies the Generation 2, Generation 3, and Generation 4 modules, as well as the Fabric switches.

Table 26: Fibre Channel Modules and Fabric Switches

Part Number	Product Name and Description
Generation 4 Modules	
DS-X9248-256K9	48-port 8-Gbps Advanced Fibre Channel switching module.
DS-X9232-256K9	32-port 8-Gbps Advanced Fibre Channel switching module.
DS-X9530-SF2A-K9	Supervisor-2A module for Cisco MDS 9500 Series switches.
DS-13SLT-FAB3	Fabric 3 module that enables the 32-port and the 48-port 8-Gbps Advanced Fibre Channel switching module to use the full 96-Gbps or 256-Gbps backplane crossbar bandwidth.
Generation 3 Modules	
DS-X9248-96K9	48-port 8-Gbps Fibre Channel switching module.
DS-X9224-96K9	24-port 8-Gbps Fibre Channel switching module.
DS-X9248-48K9	4/44-port 8-Gbps Host-Optimized Fibre Channel switching module
DS-13SLT-FAB2	Fabric 2 module that enables the 24-port and the 48-port 8-Gbps Fibre Channel switching module to use the full 96-Gbps backplane bandwidth with any-to-any connectivity.
Generation 3 Fabric Switch	
DS-C9148-K9	Cisco MDS 9148 Fabric switch. 48-port 8-Gbps Fabric switch.
Generation 2 Modules	
DS-X9148	48-port 4-Gbps Fibre Channel switching module.
DS-X9124	24-port 4-Gbps Fibre Channel switching module.
DS-X9304-18K9	18-port 4-Gbps Fibre Channel switching module with 4-Gigabit Ethernet ports.

Part Number	Product Name and Description
DS-X9112	12-port 4-Gbps Fibre Channel switching module.
DS-X9704	4-port 10-Gbps Fibre Channel switching module.
DS-X9530-SF2-K9	Supervisor-2 module for Cisco MDS 9500 Series switches.
Generation 2 Fabric Switches	
DS-C9134-K9	Cisco MDS 9134 Fabric switch. 32-port 4-Gbps Fabric switch with 2 additional 10-Gbps ports.
DS-C9124-K9	Cisco MDS 9124 Fabric switch. 24-port 4-Gbps Fabric switch.
DS-C9222i-K9	Cisco MDS 9222i Multiservice Modular switch. 18-port 4-Gbps switch with 4-Gigabit Ethernet IP storage services ports, and a modular expansion slot to host Cisco MDS 9000 Family switching and services modules.

**Note**

Generation 2 Fibre Channel switching modules are not supported on the Cisco MDS 9216 switch; however, they are supported by both the Supervisor-1 module and the Supervisor-2 module.

For detailed information about the installation and specifications for these modules and switches, refer to the hardware installation guide for your switch.

Port Groups

Each module or switch can have one or more ports in port groups that share common resources such as bandwidth and buffer credits. Port groups are defined by the hardware consisting of sequential ports. For example, ports 1 through 6, ports 7 through 12, ports 13 through 18, ports 19 through 24, ports 25 through 30, 31 through 36, and ports 37 through 42, 43 through 48 are the port groups on the 48-port 8-Gbps Advanced Fibre Channel switching modules.

[Table 27: Bandwidth and Port Groups for the Fibre Channel Modules and Fabric Switches, on page 313](#) shows the bandwidth and number of ports per port group for the Generation 2, Generation 3, and Generation 4 Fibre Channel modules, and Generation 2 and Generation 3 Fabric switches.

Table 27: Bandwidth and Port Groups for the Fibre Channel Modules and Fabric Switches

Part Number	Product Name/ Description	Number of Ports Per Port Group	Bandwidth Per Port Group (Gbps)	Maximum Bandwidth Per Port (Gbps)
Generation 4 Modules				

Part Number	Product Name/ Description	Number of Ports Per Port Group	Bandwidth Per Port Group (Gbps)	Maximum Bandwidth Per Port (Gbps)
DS-X9248-256K9	48-port 8-Gbps Advanced Fibre Channel switching module.	6	32.4 or 12.8	8 or 10 Gbps—depending on the configuration
DS-X9232-256K9	32-port 8-Gbps Advanced Fibre Channel switching module.	4	32.4 ¹⁰ or 12.8 ¹¹	8 or 10 Gbps—depending on the configuration
Generation 3 Modules				
DS-X9248-96K9	48-port 8-Gbps Fibre Channel switching module	6	12.8	8 Gbps
DS-X9224-96K9	24-port 8-Gbps Fibre Channel switching module	3	12.8	8 Gbps
DS-X9248-48K9	4/44-port 8-Gbps Host-Optimized Fibre Channel switching module	12	12.8	8/4 Gbps ¹²
Generation 3 Fabric Switches				
DS-C9148-K9 (Cisco MDS 9148 Fabric switch)	48-port 8-Gbps Fabric switch	4	32	8 Gbps
Generation 2 Modules				
DS-X9148	48-port 4-Gbps Fibre Channel switching module	12	12.8	4 Gbps
DS-X9124	24-port 4-Gbps Fibre Channel switching module	6	12.8	4 Gbps
DS-X9304-18K9 (MSM-18/4 Multiservice module)	18-port 4-Gbps Fibre Channel switching module with 4-Gigabit Ethernet ports	6	12.8	4 Gbps
DS-X9112	12-port 4-Gbps Fibre Channel switching module	3	12.8	4 Gbps
DS-X9704	4-port 10-Gbps Fibre Channel switching module	1	10	10 Gbps
Generation 2 Fabric Switches				
DS-C9134-K9 (Cisco MDS 9134 Fabric switch)	32-port 4-Gbps Fabric switch	4	16	4 Gbps
	2-port 10-Gbps Fabric switch	1	10	10 Gbps

Part Number	Product Name/ Description	Number of Ports Per Port Group	Bandwidth Per Port Group (Gbps)	Maximum Bandwidth Per Port (Gbps)
DS-C9124K9 (Cisco MDS 9124 Fabric switch)	24-port 4-Gbps Fabric switch	4	16	4 Gbps
DS-C9222i-K9 (Cisco MDS 9222i Multiservice Modular switch)	18-port 4-Gbps, 4 Gigabit Ethernet ports and a modular expansion slot.	6	12.8	4 Gbps

⁸ This bandwidth is available with the Fabric 3 module (DS-13SLT-FAB3) in the MDS 9513 switch.

⁹ This bandwidth is available with the Fabric 2 module (DS-13SLT-FAB2) in the MDS 9513 switch, and with the Supervisor-2 (DS-X9530-SF2-K9) or Supervisor-2A module (DS-X9530-SF2AK9) in the MDS 9509 switch or MDS 9506 switch.

¹⁰ [Footnote.](#)

¹¹ [Footnote.](#)

¹² A maximum of four ports (one per port group) in a 4/44-port 8-Gbps switching module can operate at 8-Gbps bandwidth in dedicated or shared mode. All the other ports can operate at a maximum of 4-Gbps in shared mode or dedicated mode.

Port Rate Modes

In Generation 2, Generation 3, and Generation 4 modules, you can configure the port rate modes. The *port rate mode* configuration is used to determine the bandwidth allocation for ports in a port group. Two port rate modes are supported:



Note In Generation 1 modules, you cannot configure the port rate modes. The mode is determined implicitly based on the port mode and line card type.



Note Port rate modes are not supported on the Cisco Fabric Switch for HP c-Class Blade System, and the Cisco Fabric Switch for IBM Blade Center.

[Table 28: Port Rate Mode Support on Generation 2, Generation 3, and Generation 4 Modules and Switches](#) shows the modules that support dedicated, shared, and the default rate modes.

Table 28: Port Rate Mode Support on Generation 2, Generation 3, and Generation 4 Modules and Switches

Part Number	Product Name/Description	Supports Dedicated Rate Mode	Supports Shared Rate Mode	Default Speed Mode and Rate Mode on All Ports
Generation 4 Modules				
DS-X9248-256K9	48-port 8-Gbps Advanced Fibre Channel switching module	Yes	Yes	Auto, Shared

Part Number	Product Name/Description	Supports Dedicated Rate Mode	Supports Shared Rate Mode	Default Speed Mode and Rate Mode on All Ports
DS-X9232-256K9	32-port 8-Gbps Advanced Fibre Channel switching module	Yes	Yes ¹⁴	Auto, Shared
Generation 3 Modules				
DS-X9248-96K9	48-Port 8-Gbps Fibre Channel switching module	Yes	Yes	Auto, Shared
DS-X9224-96K9	24-Port 8-Gbps Fibre Channel switching module	Yes	Yes ¹⁶	Auto, Shared
DS-X9248-48K9	4/44-Port 8-Gbps Host-Optimized Fibre Channel switching module	Yes	Yes ¹⁷	Auto Max 4 Gbps, Shared
Generation 3 Fabric Switches				
DS-C9148-K9 (Cisco MDS 9148 Fabric switch)	48-port 8-Gbps Fabric switch	Yes	No	Auto, Dedicated
Generation 2 Modules				
DS-X9148	48-port 4-Gbps Fibre Channel switching module ¹⁸	Yes	Yes	Auto, Shared
DS-X9124	24-port 4-Gbps Fibre Channel switching module	Yes	Yes	Auto, Shared
DS-X9304-18K9 (MSM-18/4 Multiservice module)	18-port 4-Gbps Fibre Channel switching module with 4-Gigabit Ethernet ports	Yes	Yes	Auto, Shared
DS-X9112	12-port 4-Gbps Fibre Channel switching module	Yes	No	Auto, Dedicated
DS-X9704	4-port 10-Gbps Fibre Channel switching module	Yes	No	Auto, Dedicated
Generation 2 Fabric Switches				
DS-C9134-K9 (Cisco MDS 9134 Fabric switch)	32-port 4-Gbps Fabric switch	Yes	Yes	Auto, Shared
	2-port 10-Gbps Fabric switch	Yes	No	Auto, Dedicated
DS-C9124-K9 (Cisco MDS 9124 Fabric switch)	24-port 4-Gbps Fabric switch ¹⁹	Yes	No	Auto, Dedicated

Part Number	Product Name/Description	Supports Dedicated Rate Mode	Supports Shared Rate Mode	Default Speed Mode and Rate Mode on All Ports
DS-C9222i-K9 (Cisco MDS 9222i Multiservice Modular switch)	18-port 4-Gbps Fibre Channel switch with 4-Gigabit Ethernet IP storage services ports, and a modular expansion slot to host Cisco MDS 9000 Family Switching and Services Modules	Yes	Yes	Auto, Shared

¹³ Supports shared rate mode.

¹⁴ [Footnote.](#)

¹⁵ Shared rate mode is supported on Fx ports only and no ISLs.

¹⁶ [Footnote.](#)

¹⁷ [Footnote.](#)

¹⁸ All ports in a 48-port 4-Gbps switching module can operate in dedicated rate mode with a 1-Gbps operating speed. However, if you configure one or more ports to operate in 2-Gbps or 4-Gbps dedicated rate mode, some of the other ports in the port group would have to operate in shared mode.

¹⁹ All ports in a 24-port 4-Gbps switching module can operate in dedicated rate mode with a 2-Gbps operating speed. However, if you configure one or more ports to operate in 4-Gbps dedicated rate mode, some of the other ports in the port group would have to operate in shared mode.

Dedicated Rate Mode

When port rate mode is configured as dedicated, a port is allocated required fabric bandwidth and related resources to sustain line rate traffic at the maximum operating speed configured for the port. In this mode, ports do not use local buffering and all receive buffers are allocated from a global buffer pool (see the [“Buffer Pools”](#) section on page 60-2).

Table 29: [Bandwidth Reserved for the Port Speeds on Generation 4 Switching Modules](#) , on page 317 shows the bandwidth provided by the various port speed configurations on the 8-Gbps Advanced Fibre Channel switching modules.

Table 29: Bandwidth Reserved for the Port Speeds on Generation 4 Switching Modules

Configured Speed	Reserved Bandwidth
Auto	8 Gbps
8-Gbps	
Auto with 4-Gbps maximum	4 Gbps
4-Gbps	
Auto with 2-Gbps maximum	2 Gbps
2-Gbps	
1-Gbps	1 Gbps

[Table 30: Bandwidth Reserved for the Port Speeds on Generation 3 Switching Modules](#), on page 318 shows the bandwidth provided by the various port speed configurations on the 8-Gbps Fibre Channel switching modules.

Table 30: Bandwidth Reserved for the Port Speeds on Generation 3 Switching Modules

Configured Speed	Reserved Bandwidth
Auto	8 Gbps
8-Gbps	
Auto with 4-Gbps maximum	4 Gbps
4-Gbps	
Auto with 2-Gbps maximum	2 Gbps
2-Gbps	
1-Gbps	1 Gbps

[Table 31: Bandwidth Reserved for the Port Speeds on Generation 2 Switching Modules](#), on page 318 shows the amount of bandwidth reserved for a configured port speed on 4-Gbps switching modules.

Table 31: Bandwidth Reserved for the Port Speeds on Generation 2 Switching Modules

Configured Speed	Reserved Bandwidth
Auto	4 Gbps
4-Gbps	
Auto with 2-Gbps maximum	2 Gbps
2-Gbps	
1-Gbps	1 Gbps



Note The 4-Port 10-Gbps Fibre Channel module ports in auto mode only support auto speed mode at 10 Gbps.

Shared Rate Mode

When port rate mode is configured as shared, multiple ports within a port group share data paths to the switch fabric so that fabric bandwidth and related resources are shared. Often, the available bandwidth to the switch fabric may be less than the negotiated operating speed of a port. Ports in this mode use local buffering for the BB_credit buffers.

All ports in 8-Gbps Advanced Fibre Channel switching modules where bandwidth is shared support 1-Gbps, 2-Gbps, 4-Gbps, or 8 Gbps traffic. However, it is possible to configure one or more ports in a port group to operate in dedicated rate mode with 1-Gbps, 2-Gbps, 4-Gbps, or 8 Gbps operating speed.

All ports in 4-Gbps Fibre Channel switching modules where bandwidth is shared support 1-Gbps, 2-Gbps, or 4-Gbps traffic. However, it is possible to configure one or more ports in a port group to operate in dedicated rate mode with 1-Gbps, 2-Gbps, or 4-Gbps operating speed.

All ports in the 32-Port or 48-Port 8-Gbps Advanced Fibre Channel modules where bandwidth is shared support 1-Gbps, 2-Gbps, 4-Gbps, or 8-Gbps traffic in a maximum of 32 or 48 ports.

All ports in the 48-Port and 24-Port 8-Gbps Fibre Channel switching modules where bandwidth is shared support 1-Gbps, 2-Gbps, 4-Gbps, or 8-Gbps traffic.

In the 4/44-Port 8-Gbps Host-Optimized Fibre Channel switching module, all the ports where bandwidth is shared support 1-Gbps, 2-Gbps, 4-Gbps in a maximum of 44 ports, or 8 Gbps in a maximum of 4 ports.

Dedicated Rate Mode Configurations for the 8-Gbps Modules

Table 32: Dedicated Rate Mode Bandwidth Reservation for Generation 4 Fibre Channel Modules , on page 319 shows the maximum possible dedicated rate mode configuration scenarios for the Generation 4 Fibre Channel modules.

Table 32: Dedicated Rate Mode Bandwidth Reservation for Generation 4 Fibre Channel Modules

Part Number	Product Name/Description	Dedicated Bandwidth per Port	Maximum Allowed Ports That Can Come Up	Ports in Shared Mode
DS-X9248-256K9	48-port 8-Gbps Advanced Fibre Channel switching module	10 Gbps	24 Ports	All the remaining ports are 8 Gbps shared.
		8 Gbps	32 Ports	
		4 Gbps	48 Ports	
		2 Gbps	48 Ports	
		1 Gbps	48 Ports	
DS-X9232-256K9	32-port 8-Gbps Advanced Fibre Channel switching module	10 Gbps	24 Ports	All the remaining ports are 8 Gbps shared.
		8 Gbps	32 Ports	
		4 Gbps	32 Ports	
		2 Gbps	32 Ports	
		1 Gbps	32 Ports	

Table 33: Dedicated Rate Mode Bandwidth Reservation for Generation 3 Fibre Channel Modules , on page 319 shows the maximum possible dedicated rate mode configuration scenarios for the Generation 3 Fibre Channel modules.

Table 33: Dedicated Rate Mode Bandwidth Reservation for Generation 3 Fibre Channel Modules

Part Number	Product Name/Description	Dedicated Bandwidth per Port	Maximum Allowed Ports That Can Come Up	Ports in Shared Mode
DS-X9224-96K9	24-port 8-Gbps Fibre Channel switching module	8 Gbps	8 Ports	All the remaining ports are 8 Gbps shared.
		4 Gbps	24 Ports	

Part Number	Product Name/ Description	Dedicated Bandwidth per Port	Maximum Allowed Ports That Can Come Up	Ports in Shared Mode
DS-X9248-96K9	48-port 8-Gbps Fibre Channel switching module	8 Gbps	8 Ports	All the remaining ports are 8 Gbps shared.
		4 Gbps	24 Ports	
		2 Gbps	48 Ports	
DS-X9248-48K9	4/44-port 8-Gbps Host-Optimized Fibre Channel switching module	8 Gbps	4 Ports	All the remaining ports are 4 Gbps shared (8 Gbps of bandwidth can be provided only to one port per port group in dedicated or shared rate mode).
		4 Gbps	12 Ports	
		2 Gbps	24 Ports	
		1 Gbps	48 Ports	

Port Speed

The port speed on an interface, combined with the rate mode, determines the amount of shared resources available to the ports in the port group on a 48-port, 24-port 4-Gbps, or any 8-Gbps Fibre Channel switching module. Especially in the case of dedicated rate mode, the port group resources are reserved even though the bandwidth is not used. For example, on Generation 2 modules, if an interface is configured for autosensing (auto) and dedicated rate mode, then 4 Gbps of bandwidth is reserved even though the maximum operating speed is 2 Gbps. For the same interface, if autosensing with a maximum speed of 2 Gbps (auto max 2000) is configured, then only 2 Gbps of bandwidth is reserved and the unused 2 Gbps is shared with the other interface in the port group.



Note

The Generation 2, 4-port 10-Gbps switching module supports 10-Gbps traffic only.

- On Generation 2, 4-Gbps modules, setting the port speed to auto enables autosensing, which negotiates to a maximum speed of 4 Gbps.
- On Generation 3, 8-Gbps modules, setting the port speed to auto enables autosensing, which negotiates to a maximum speed of 8 Gbps.
- On Generation 4, 8-Gbps modules, setting the port speed to auto enables autosensing, which negotiates to a maximum speed of 8 Gbps.

Dynamic Bandwidth Management

On port switching modules where bandwidth is shared, the bandwidth available to each port within a port group can be configured based on the port rate mode and speed configurations. Within a port group, some ports can be configured in dedicated rate mode while others operate in shared mode.

Ports configured in dedicated rate mode are allocated the required bandwidth to sustain a line rate of traffic at the maximum configured operating speed, and ports configured in shared mode share the available remaining bandwidth within the port group. Bandwidth allocation among the shared mode ports is based on the operational speed of the ports. For example, if four ports operating at speeds 1 Gbps, 1 Gbps, 2 Gbps, and 4 Gbps share bandwidth of 8 Gbps, the ratio of allocation would be 1:1:2:4.

Unutilized bandwidth from the dedicated ports is shared among only the shared ports in a port group as per the ratio of the configured operating speed. A port cannot be brought up unless the reserved bandwidth is guaranteed for the shared ports. For dedicated ports, configured bandwidth is taken into consideration while calculating available bandwidth for the port group. This behavior can be changed using bandwidth fairness by using the **rate-mode bandwidth-fairness module** *number* command.

For example, consider a 48-port 8-Gbps module. This module has 6 ports per port group with 12.8 Gbps bandwidth. Ports 3 to 6 are configured at 4 Gbps. If the first port is configured at 8 Gbps dedicated rate mode, and the second port is configured at 4-Gbps dedicated rate mode, then no other ports can be configured at 4 Gbps or 8 Gbps because the left over bandwidth of 0.8 Gbps ($12.8 - (8 + 4)$) cannot meet the required 0.96 Gbps for the remaining four ports. A minimum of 0.24 Gbps reserved bandwidth is required for the for the rest of the four ports. However, if the two ports (for example, 5 and 6) are taken out of service (which is not same as shutdown), required reserved bandwidth for the two ports (3 and 4) is 0.48 and port 2 can be configured at 4 Gbps in dedicated rate mode. This behavior can be overridden by the bandwidth fairness command in which case reserved bandwidth is not enforced. Once the port is up, ports 3 and 4 can share the unutilized bandwidth from ports 1 and 2.

Bandwidth Reservation: 48-Port 96-Gbps Fibre Channel Module

RateMode Configuration Macro	Description
Dedicated 4 Gbps on the first port of each group and the remaining ports 8 Gbps shared	Allocates a rate mode of 4 Gbps on the first port of each group and the remaining ports share 8 Gbps depending on the operational speed of the ports.
Dedicated 8 Gbps on the first port of each group and the remaining ports 8 Gbps shared	Allocates a rate mode of 8 Gbps on the first port of each group and the remaining ports share 8 Gbps depending on the operational speed of the ports.
Shared 8 Gbps on all ports (initial and default settings)	Allocates a rate mode of 8 Gbps on all the available ports. This is the default setting.

Bandwidth Reservation: 48-Port 48-Gbps Fibre Channel Module

RateMode Configuration Macro	Description
Dedicated 2 Gbps on the first port of each group and the remaining ports 4 Gbps shared	Allocates a rate mode of 2 Gbps on the first port of each group and the remaining ports share 4 Gbps depending on the operational speed of the ports.
Dedicated 8 Gbps on the first port of each group and the remaining ports 4 Gbps shared	Allocates a rate mode of 8 Gbps on the first port of each group and the remaining ports share 4 Gbps depending on the operational speed of the ports.
Shared auto with maximum of 4 Gbps on all ports (initial and default settings)	Allocates a maximum rate mode of 4 Gbps on all the available ports. This is the default setting.

Bandwidth Reservation: 24-Port 48-Gbps Fibre Channel Module

RateMode Configuration Macro	Description
Dedicated 8 Gbps on the first port of each group and the remaining ports 8 Gbps shared	Allocates a rate mode of 8Gbps on the first port of each group and the remaining ports share 8 Gbps depending on the operational speed of the ports.
Shared Auto on all ports (initial and default settings)	Allocates a rate mode of 8 Gbps on all the available ports. This is the default setting.

Bandwidth Reservation: 48-Port 256-Gbps Advanced Fibre Channel Module

RateMode Configuration Macro	Description
Dedicated 8 Gbps on the first 4 ports in each 6-port port group and the remaining ports 8 Gbps shared	Allocates a rate mode of 8 Gbps on the first 4 ports in each 6-port port group and the remaining ports share 8 Gbps depending on the operational speed of the ports.
Dedicated 8 Gbps on the first port of each group and the remaining ports 8 Gbps shared	Allocates a rate mode of 8 Gbps on the first port of each group and the remaining ports share 8 Gbps depending on the operational speed of the ports.
Shared 8 Gbps on all ports	Allocates a rate mode of 8 Gbps on all the available ports. This is the default setting.
Dedicated 4 Gbps on all ports	Allocates a rate mode of 4 Gbps on all the available ports.
Dedicated 10 Gbps on following ports: <ul style="list-style-type: none"> • 4, 5, 6, 7, 8, 10 (ports 1,2, 3, 9, 11, 12 disabled) • 16, 17, 18, 19, 20, 22 (ports 13, 14, 15, 21, 23, 24 disabled) • 28, 29, 30, 31, 32, 34 (ports 25, 26, 27, 33, 35, 36 disabled) • 40, 41, 42, 43, 44, 46 (ports 37, 38, 39, 45, 47, 48 disabled) 	Allocates a rate mode of 10 Gbps on all the available ports.

Bandwidth Reservation: 32-Port 256-Gbps Advanced Fibre Channel Module

RateMode Configuration Macro	Description
Dedicated 8 Gbps on all ports—initial and default settings	Allocates a rate mode of 8 Gbps on all the available ports.
Shared 8 Gbps on all ports—initial and default settings	Allocates a rate mode of shared 8 Gbps on all the available ports.

RateMode Configuration Macro	Description
Dedicated 10 Gbps on following ports: <ul style="list-style-type: none"> • 2, 3, 4, 5, 6, 8 (ports 1 and 7 disabled) • 10, 11, 12, 13, 14, 16 (ports 9 and 15 disabled) • 18, 19, 20, 21, 22, 24 (ports 17 and 23 disabled) • 26, 27, 28, 29, 30, 32 (ports 25 and 31 disabled) 	Allocates a rate mode of 10Gbps on the following ports.

Out-of-Service Interfaces

On supported modules and fabric switches, you might need to allocate all the shared resources for one or more interfaces to another interface in the port group or module. You can take interfaces out of service to release shared resources that are needed for dedicated bandwidth. When an interface is taken out of service, all shared resources are released and made available to the other interface in the port group or module. These shared resources include bandwidth for the shared mode port, rate mode, BB_credits, and extended BB_credits. All shared resource configurations are returned to their default values when the interface is brought back into service. Corresponding resources must be made available in order for the port to be successfully returned to service.



Caution

If you need to bring an interface back into service, you might disrupt traffic if you need to release shared resources from other interfaces in the same port group.

Oversubscription Ratio Restrictions

The 48-port and 24-port 4-Gbps, and all 8-Gbps Fibre Channel switching modules support oversubscription on switches with shared rate mode configurations. By default, all 48-port and 24-port 4-Gbps, and 8-Gbps Fibre Channel switching modules have restrictions on oversubscription ratios enabled. As of Cisco SAN-OS Release 3.1(1) and NX-OS Release 4.1(1), you can disable restrictions on oversubscription ratios.

[Table 34: Bandwidth Allocation for Oversubscribed Interfaces, on page 323](#) describes the bandwidth allocation for oversubscribed interfaces configured in shared mode on the 4-Gbps and 8-Gbps modules.

Table 34: Bandwidth Allocation for Oversubscribed Interfaces

Switching Module	Configured Speed	Reserved Bandwidth (Gbps)		Maximum Bandwidth (Gbps)
		Ratios enabled	Ratios disabled	
24-Port 8-Gbps Fibre Channel Module	Auto 8 Gbps	0.8	0.8	8
	Auto Max 4 Gbps	0.4	0.4	4
	Auto Max 2 Gbps	0.2	0.2	2
4/44-Port 8-Gbps Host-Optimized Fibre Channel Module	8 Gbps	0.87	0.16	8
	Auto Max 4 Gbps	0.436	0.08	4
	Auto Max 2 Gbps	0.218	0.04	2
	1 Gbps	0.109	0.02	1

Switching Module	Configured Speed	Reserved Bandwidth (Gbps)		Maximum Bandwidth (Gbps)
		Ratios enabled	Ratios disabled	
48-port 4-Gbps Fibre Channel switching module	Auto 4 Gbps	0.8	0.09	4
	Auto Max 2 Gbps	0.4	0.045	2
	1 Gbps	0.2	0.0225	1
24-port 4-Gbps Fibre Channel switching module	Auto 4 Gbps	1	0.27	4
	Auto Max 2 Gbps	0.5	0.135	2
	1 Gbps	0.25	0.067	1

All ports in the 48-port and 24-port 4-Gbps modules can be configured to operate at 4 Gbps in shared mode even if other ports in the port group are configured in dedicated mode, regardless of available bandwidth. However, when oversubscription ratio restrictions are enabled, you may not have all shared 4-Gbps module ports operating at 4 Gbps.

All ports in the 48-port, 32-Port, and 24-port 8-Gbps modules can be configured to operate at 8 Gbps in shared mode even if other ports in the port group are configured in dedicated mode, regardless of available bandwidth. However, when oversubscription ratio restrictions are enabled you may not have all shared 8-Gbps module ports operating at 8 Gbps.

On the 48-port, 32-Port, and 24-port 8-Gbps modules, if you have configured one 8-Gbps dedicated port in one port group, no other ports in the same port group can be configured to operate at 8-Gbps dedicated mode. You can have any number of 8-Gbps shared and 4-Gbps dedicated or shared ports. On the 4/44-port 8-Gbps module, only one port per port group can be configured in 8-Gbps dedicated or shared mode.

In the following example, a 24-port 4-Gbps module has oversubscription ratios enabled and three dedicated ports in one port group operating at 4-Gbps. No other ports in the same port group can be configured to operate at 4 Gbps.

```
switch# show port-resources module 8
Module 8
  Available dedicated buffers are 5478

Port-Group 1
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 0.8 Gbps
  Allocated dedicated bandwidth is 12.0 Gbps
-----
Interfaces in the Port-Group      B2B Credit   Bandwidth   Rate Mode
                                Buffers      (Gbps)
-----
fc8/1                            16           4.0         dedicated
fc8/2                            16           4.0         dedicated
fc8/3                            16           4.0         dedicated
fc8/4 (out-of-service)
fc8/5 (out-of-service)
fc8/6 (out-of-service)
```

For dedicated ports, oversubscription ratio restrictions do not apply to the shared pool in port groups. So if oversubscription ratio restrictions are disabled, and you have configured three 4-Gbps dedicated ports in one port group, then you can configure all other ports in the same port group to operate at a shared rate of 4 Gbps.

In the following example, a 48-port module has a group of six ports, four dedicated ports are operating at 8 Gbps, and the two shared ports are also operating at 8 Gbps:

```

switch# show port-resources module 5
Module 5
Available dedicated buffers for global buffer #0 [port-group 1] are 3970
  Available dedicated buffers for global buffer #1 [port-group 2] are 3970
  Available dedicated buffers for global buffer #2 [port-group 3] are 3970
  Available dedicated buffers for global buffer #3 [port-group 4] are 3970
  Available dedicated buffers for global buffer #4 [port-group 5] are 3058
  Available dedicated buffers for global buffer #5 [port-group 6] are 3058
  Available dedicated buffers for global buffer #6 [port-group 7] are 3970
  Available dedicated buffers for global buffer #7 [port-group 8] are 3970
Port-Group 1
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
-----
  Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                   Buffers      (Gbps)
-----
  fc5/1                             32           8.0    shared
  fc5/2                             32           8.0    shared
  fc5/3                             32           8.0    shared
  fc5/4                             32           8.0    shared
Port-Group 2
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
-----
  Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                   Buffers      (Gbps)
-----
  fc5/5                             32           8.0    shared
  fc5/6                             32           8.0    shared
  fc5/7                             32           8.0    shared
  fc5/8                             32           8.0    shared
Port-Group 3
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
-----
  Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                   Buffers      (Gbps)
-----
  fc5/9                             32           8.0    shared
  fc5/10                            32           8.0    shared
  fc5/11                            32           8.0    shared
  fc5/12                            32           8.0    shared
Port-Group 4
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
-----
  Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                   Buffers      (Gbps)
-----
  fc5/13                            32           8.0    shared
  fc5/14                            32           8.0    shared
  fc5/15                            32           8.0    shared
  fc5/16                            32           8.0    shared
Port-Group 5
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 16.4 Gbps
  Allocated dedicated bandwidth is 16.0 Gbps
-----

```

Oversubscription Ratio Restrictions

```

Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                Buffers          (Gbps)
-----
fc5/17                          32          8.0    shared
fc5/18                          32          8.0    shared
fc5/19                          500         8.0    dedicated
fc5/20                          500         8.0    dedicated
Port-Group 6
Total bandwidth is 32.4 Gbps
Total shared bandwidth is 16.4 Gbps
Allocated dedicated bandwidth is 16.0 Gbps
-----
Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                Buffers          (Gbps)
-----
fc5/21                          500         8.0    dedicated
fc5/22                          500         8.0    dedicated
fc5/23                          32          8.0    shared
fc5/24                          32          8.0    shared
Port-Group 7
Total bandwidth is 32.4 Gbps
Total shared bandwidth is 32.4 Gbps
Allocated dedicated bandwidth is 0.0 Gbps
-----
Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                Buffers          (Gbps)
-----
fc5/25                          32          8.0    shared
fc5/26                          32          8.0    shared
fc5/27                          32          8.0    shared
fc5/28                          32          8.0    shared
Port-Group 8
Total bandwidth is 32.4 Gbps
Total shared bandwidth is 32.4 Gbps
Allocated dedicated bandwidth is 0.0 Gbps
-----
Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                Buffers          (Gbps)
-----
fc5/29                          32          8.0    shared
fc5/30                          32          8.0    shared
fc5/31                          32          8.0    shared
fc5/32                          32          8.0    shared
Isola-13# show port-resources module 13
Module 13
Available dedicated buffers for global buffer #0 [port-group 1] are 3880
Available dedicated buffers for global buffer #1 [port-group 2] are 3880
Available dedicated buffers for global buffer #2 [port-group 3] are 3880
Available dedicated buffers for global buffer #3 [port-group 4] are 3056
Available dedicated buffers for global buffer #4 [port-group 5] are 3880
Available dedicated buffers for global buffer #5 [port-group 6] are 3880
Available dedicated buffers for global buffer #6 [port-group 7] are 3880
Available dedicated buffers for global buffer #7 [port-group 8] are 3880
Port-Group 1
Total bandwidth is 32.4 Gbps
Total shared bandwidth is 32.4 Gbps
Allocated dedicated bandwidth is 0.0 Gbps
-----
Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                Buffers          (Gbps)
-----
fc13/1                          32          8.0    shared
fc13/2                          32          8.0    shared
fc13/3                          32          8.0    shared

```



```

fc13/4                32      8.0  shared
fc13/5                32      8.0  shared
fc13/6                32      8.0  shared

```

Port-Group 2

```

Total bandwidth is 32.4 Gbps
Total shared bandwidth is 32.4 Gbps
Allocated dedicated bandwidth is 0.0 Gbps

```

Interfaces in the Port-Group	B2B Credit Buffers	Bandwidth (Gbps)	Rate Mode
fc13/7	32	8.0	shared
fc13/8	32	8.0	shared
fc13/9	32	8.0	shared
fc13/10	32	8.0	shared
fc13/11	32	8.0	shared
fc13/12	32	8.0	shared

Port-Group 3

```

Total bandwidth is 32.4 Gbps
Total shared bandwidth is 32.4 Gbps
Allocated dedicated bandwidth is 0.0 Gbps

```

Interfaces in the Port-Group	B2B Credit Buffers	Bandwidth (Gbps)	Rate Mode
fc13/13	32	8.0	shared
fc13/14	32	8.0	shared
fc13/15	32	8.0	shared
fc13/16	32	8.0	shared
fc13/17	32	8.0	shared
fc13/18	32	8.0	shared

Port-Group 4

```

Total bandwidth is 32.4 Gbps
Total shared bandwidth is 0.4 Gbps
Allocated dedicated bandwidth is 32.0 Gbps

```

Interfaces in the Port-Group	B2B Credit Buffers	Bandwidth (Gbps)	Rate Mode
fc13/19	250	8.0	dedicated
fc13/20	250	8.0	dedicated
fc13/21	250	8.0	dedicated
fc13/22	250	8.0	dedicated
fc13/23	32	8.0	shared
fc13/24	32	8.0	shared

Port-Group 5

```

Total bandwidth is 32.4 Gbps
Total shared bandwidth is 32.4 Gbps
Allocated dedicated bandwidth is 0.0 Gbps

```

Interfaces in the Port-Group	B2B Credit Buffers	Bandwidth (Gbps)	Rate Mode
fc13/25	32	8.0	shared
fc13/26	32	8.0	shared
fc13/27	32	8.0	shared
fc13/28	32	8.0	shared
fc13/29	32	8.0	shared
fc13/30	32	8.0	shared

Port-Group 6

```

Total bandwidth is 32.4 Gbps
Total shared bandwidth is 32.4 Gbps
Allocated dedicated bandwidth is 0.0 Gbps

```

```

Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                Buffers          (Gbps)
-----
fc13/31                          32          8.0    shared
fc13/32                          32          8.0    shared
fc13/33                          32          8.0    shared
fc13/34                          32          8.0    shared
fc13/35                          32          8.0    shared
fc13/36                          32          8.0    shared
Port-Group 7
Total bandwidth is 32.4 Gbps
Total shared bandwidth is 32.4 Gbps
Allocated dedicated bandwidth is 0.0 Gbps
-----
Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                Buffers          (Gbps)
-----
fc13/37                          32          8.0    shared
fc13/38                          32          8.0    shared
fc13/39                          32          8.0    shared
fc13/40                          32          8.0    shared
fc13/41                          32          8.0    shared
fc13/42                          32          8.0    shared
Port-Group 8
Total bandwidth is 32.4 Gbps
Total shared bandwidth is 32.4 Gbps
Allocated dedicated bandwidth is 0.0 Gbps
-----
Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                Buffers          (Gbps)
-----
fc13/43                          32          8.0    shared
fc13/44                          32          8.0    shared
fc13/45                          32          8.0    shared
fc13/46                          32          8.0    shared
fc13/47                          32          8.0    shared
fc13/48                          32          8.0    shared
...

```

When disabling restrictions on oversubscription ratios, all ports in shared mode on 48-port and 24-port 4-Gbps or any 8-Gbps Fibre Channel switching modules must be shut down. When applying restrictions on oversubscription ratios, you must take shared ports out of service.

**Note**

When restrictions on oversubscription ratios are disabled, the bandwidth allocation among the shared ports is proportionate to the configured speed. If the configured speed is auto on Generation 2 modules, then bandwidth is allocated assuming a speed of 4 Gbps. For example, if you have three shared ports configured at 1, 2, and 4 Gbps, then the allocated bandwidth ratio is 1:2:4. As of Cisco SAN-OS Release 3.0 and NX-OS Release 4.1(1) or when restrictions on oversubscription ratios are enabled, the port bandwidths are allocated in equal proportions, regardless of port speed, so, the bandwidth allocation for the same three ports mentioned in the example would be 1:1:1.

Bandwidth Fairness

This feature improves fairness of bandwidth allocation among all ports and provides better throughput average to individual data streams. Bandwidth fairness can be configured per module.

As of Cisco SAN-OS Release 3.1(2), all 48-port and 24-port 4-Gbps Fibre Channel switching modules, as well as 18-port Fibre Channel/4-port Gigabit Ethernet Multiservice modules, have bandwidth fairness enabled

by default. As of Cisco NX-OS Release 4.1(1), all the 8-Gbps Fibre Channel switching modules have bandwidth fairness enabled by default.

**Caution**

When you disable or enable bandwidth fairness, the change does not take effect until you reload the module.

Use the `show module bandwidth-fairness` command to check whether ports in a module are operating with bandwidth fairness enabled or disabled.

```
switch# show module 2 bandwidth-fairness
Module 2 bandwidth-fairness is enabled
```

**Note**

This feature is supported only on the 48-port and 24-port 4-Gbps modules, the 8-Gbps modules, and the 18/4-port Multiservice Module (MSM).

Upgrade or Downgrade Scenario

When you are upgrading from a release earlier than Cisco SAN-OS Release 3.1(2), all modules operate with bandwidth fairness disabled until the next module reload. After the upgrade, any new module that is inserted has bandwidth fairness enabled.

When you are downgrading to a release earlier than Cisco SAN-OS Release 3.1(2), all modules keep operating in the same bandwidth fairness configuration prior to the downgrade. After the downgrade, any new module that is inserted has bandwidth fairness disabled.

**Note**

After the downgrade, any insertion of a module or module reload will have bandwidth fairness disabled.

Guidelines and Limitations

This section includes the following topics:

Combining Generation 1, Generation 2, Generation 3, and Generation 4 Modules

Cisco MDS NX-OS Release 6.x and later supports combining Generation 1, Generation 2, Generation 3, and Generation 4 modules and switches with the following considerations:

- MDS NX-OS Release 4.1(1) and later features are not supported on the following Generation 1 switches and modules:
 - Supervisor 1 module
 - 4-Port IP Storage Services module
 - 8-Port IP Storage Services module
 - MDS 9216 switch
 - MDS 9216A switch
 - MDS 9020 switch
 - MDS 9120 switch
 - MDS 9140 switch

- Supervisor-1 modules must be upgraded to Supervisor-2 modules on the MDS 9506 and MDS 9509 Directors.
- IPS-4 and IPS-8 modules must be upgraded to the MSM-18/4 Multiservice modules.
- Fabric 1 modules must be upgraded to Fabric 2 modules on the MDS 9513 Director to use the 48-port or the 24-port 8-Gbps module.
- Fabric 2 modules must be upgraded to Fabric 3 modules on the MDS 9513 Director to get the maximum backplane bandwidth of 256 Gbps.
- Cisco Fabric Manager Release 4.x supports MDS SAN-OS Release 3.x and NX-OS 4.x in mixed mode through Interswitch Link (ISL) connectivity.



Note When a Cisco or another vendor switch port is connected to a Generation 1 module port (ISL connection), the receive buffer-to-buffer credits of the port connected to the Generation 1 module port must not exceed 255.

Local Switching Limitations

All ports in the module must be in shared mode.

- Use the **switchport ratemode shared** command to ensure that all the ports in the module are in shared mode.
- No E ports are allowed in the module because E ports must be in dedicated mode.

Port Index Limitations

Cisco MDS 9000 switches allocate index identifiers for the ports on the modules. These port indexes cannot be configured. You can combine Generation 1, Generation 2, Generation 3, and Generation 4 switching modules, with either Supervisor-1 modules or Supervisor-2 modules. However, combining switching modules and supervisor modules has the following port index limitations:

- Supervisor-1 modules only support a maximum of 252 port indexes, regardless of the type of switching modules.
- Supervisor-2 modules support a maximum of 1020 port indexes when all switching modules in the chassis are Generation 2 or Generation 3.
- Supervisor-2 modules only support a maximum of 252 port indexes when only Generation 1 switching modules, or a combination of Generation 1, Generation 2, Generation 3, or Generation 4 switching modules are installed in the chassis.



Note On a switch with the maximum limit of 252 as port index, any new module that exceeds the limit does not power up when installed.

You can use the **show port index-allocation** command to display the allocation of port indexes on the switch.

```
switch# show port index-allocation
Module index distribution:
-----+
Slot | Allowed |      Allotted indices info      |
    | range  | Total |      Index values      |
----|-----|-----|-----|
1 | ----- | - | (None) |
```

2		-----		-		(None)	
3		-----		-		(None)	
4		-----		-		(None)	
5		0-1023		32		0-31	
6		-----		-		(None)	
9		-----		-		(None)	
10		-----		-		(None)	
11		-----		-		(None)	
12		-----		-		(None)	
13		0-1023		48		32-79	
SUP		253-255		3		253-255	

Generation 1 switching modules have specific numbering requirements. If these requirements are not met, the module does not power up. The port index numbering requirements include the following:

- If port indexes in the range of 256 to 1020 are assigned to operational ports, Generation 1 switching modules do not power up.
- A block of contiguous port indexes is available. If this block of port indexes is not available, Generation 1 modules do not power up. [Table 35: Port Index Requirements for Generation 1 Modules](#), on page 331 shows the port index requirements for the Generation 1 modules.



Note If the switch has Supervisor-1 modules, the block of 32 contiguous port indexes must begin on the slot boundary. The slot boundary for slot 1 is 0, for slot 2 is 32, and so on. For Supervisor-2 modules, the contiguous block can start anywhere.

Table 35: Port Index Requirements for Generation 1 Modules

Generation 1 Module	Number of Port Indexes Required	
	Supervisor-1 Module	Supervisor-2 Module
16-port 2-Gbps Fibre Channel module	16	16
32-port 2-Gbps Fibre Channel module	32	32
8-port Gigabit Ethernet IP Storage Services module	32	32
4-port Gigabit Ethernet IP Storage Services module	32	16
32-port 2-Gbps Fibre Channel Storage Services Module (SSM).	32	32
14-port Fibre Channel/2-port Gigabit Ethernet Multiprotocol Services (MPS-14/2) module	32	22

The allowed mix of Generation 1 and Generation 2 switching modules in a chassis is determined at run-time, either when booting up the switch or when installing the modules. In some cases, the sequence in which switching modules are inserted into the chassis determines if one or more modules is powered up.

When a module does not power up because of a resource limitation, you can display the reason by using the **show module** command.

When a module does not power up because of a resource limitation, you can see the reason by viewing the module information in the Information pane.

```
switch# show module
```

```

Mod  Ports  Module-Type                      Model                      Status
---  ---
5    32      1/2/4/8/10 Gbps Advanced FC Module DS-X9232-256K9           ok
7    0      Supervisor/Fabric-2                DS-X9530-SF2-K9         active *
13   48      1/2/4/8/10 Gbps Advanced FC Module DS-X9248-256K9           ok
Mod  Sw      Hw      World-Wide-Name(s) (WWN)
---  ---
5    5.2(2)    0.207   21:01:00:0d:ec:b7:28:c0 to 21:20:00:0d:ec:b7:28:c0
7    5.2(2)    1.9     --
13   5.2(2)    0.212   23:01:00:0d:ec:b7:28:c0 to 23:30:00:0d:ec:b7:28:c0
Mod  MAC-Address(es)                  Serial-Num
---  ---
5    68-ef-bd-a8-45-cc to 68-ef-bd-a8-45-d0 JAF1450CHQT
7    00-24-c4-60-00-f8 to 00-24-c4-60-00-fc JAE141502L2
13   68-ef-bd-a8-40-00 to 68-ef-bd-a8-40-04 JAF1450BMBP
Xbar Ports  Module-Type                      Model                      Status
---  ---
1    0      Fabric Module 3                  DS-13SLT-FAB3            ok
2    0      Fabric Module 3                  DS-13SLT-FAB3            ok
Xbar Sw      Hw      World-Wide-Name(s) (WWN)
---  ---
1    NA      0.4     --
2    NA      0.4     --
Xbar MAC-Address(es)                  Serial-Num
---  ---
1    NA      JAF1451AMHG
2    NA      JAF1451AMHN
* this terminal session

```

The running configuration is updated when modules are installed. If you save the running configuration to the startup configuration (using the `copy running-config startup-config` command), during reboot the switch powers up the same set of modules as before the reboot regardless of the sequence in which the modules initialize. You can use the **show port index-allocation startup** command to display the index allocation the switch uses at startup.

```

switch# show port index-allocation startup
Startup module index distribution:
-----+
Slot | Allowed |      Alloted indices info      |
    | range  | Total |      Index values      |
-----+-----+-----+
1    | ----- | 34    | 0-31,80-81            |
2    | ----- | 32    | 32-63                  |
3    | ----- | 16    | 64-79                  |
4    | ----- | 48    | 96-127,224-239        |
SUP  | 253-255 | 3      | 253-255                |

```

**Note**

The output of the **show port index-allocation startup** command does not display anything in the Allowed range column because the command extracts the indices from the persistent storage service (PSS) and displaying an allowed range for startup indices is meaningless.

If a module fails to power up, you can use the **show module slot recovery-steps** command to display the reason.

For information on recovering a module powered-down because port indexes are not available, refer to the *Cisco MDS 9000 Family Troubleshooting Guide*.



Tip Whenever using mixed Generation 1 and Generation 2 modules, power up the Generation 1 modules first. During a reboot of the entire switch, the Generation 1 modules power up first (default behavior).

PortChannel Limitations

PortChannels have the following restrictions:

- The maximum number of PortChannels allowed is 256 if all switching modules are Generation 2 or Generation 3, or both.
- The maximum number of PortChannels allowed is 128 whenever there is a Generation 1 switching module in use with a Generation 2 or Generation 3 switching module.
- Ports need to be configured in dedicated rate mode on the Generation 2 and Generation 3 switching module interfaces to be used in the PortChannel.



Note The number of PortChannels allowed does not depend on the type of supervisor module. However, Generation 3 modules require the Supervisor 2 module on the MDS 9506 and 9509 switches.

The Generation 1, Generation 2, and Generation 3 modules have the following restrictions for PortChannel configuration:

- Generation 1 switching module interfaces do not support auto speed with a maximum of 2 Gbps.
- Generation 1 and Generation 2 module interfaces do not support auto speed with maximum of 4 Gbps.
- Generation 2 and Generation 3 switching module interfaces cannot be forcefully added to a PortChannel if sufficient resources are not available.



Note Before adding a Generation 2 or Generation 3 interface to a PortChannel, use the **show port-resources module** command to check for resource availability.

When configuring PortChannels on switches with Generation 1, Generation 2, and Generation 3 switching modules, follow one of these procedures:

- Configure the PortChannel, and then configure the Generation 2 and Generation 3 interfaces to auto with a maximum of 2 Gbps.
- Configure the Generation 1 switching modules followed by the Generation 2 switching modules, and then the Generation 3 switching modules, and then configure the PortChannel.

When configuring PortChannels on switches with only Generation 2 and Generation 3 switching modules, follow one of these procedures:

- Configure the PortChannel, and then configure the Generation 3 interfaces to auto with a maximum of 4 Gbps.
- Configure the Generation 2 switching modules, followed by the Generation 3 switching modules, and then configure the PortChannel.

[Table 36: PortChannel Configuration and Addition Results , on page 334](#) describes the results of adding a member to a PortChannel for various configurations.

Table 36: PortChannel Configuration and Addition Results

PortChannel Members	Configured Speed		New Member Type	Addition Type	Result
	PortChannel	New Member			
No members	Any	Any	Generation 1 or Generation 2 or Generation 3 or Generation 4	Force	Pass
	Auto	Auto	Generation 1 or Generation 2 or Generation 3 or Generation 4	Normal or force	Pass
	Auto	Auto max 2000	Generation 2 or Generation 3 or Generation 4	Normal	Fail
				Force	Pass or fail ²⁰
	Auto	Auto max 4000	Generation 3 or Generation 4		
	Auto max 2000	Auto	Generation 2 or Generation 3 or Generation 4	Normal	Fail
				Force	Pass
	Auto max 2000	Auto max 4000	Generation 3 or or Generation 4		
	Auto max 4000	Auto	Generation 2 or Generation 3 or or Generation 4		
Generation 1 interfaces	Auto	Auto	Generation 2 or Generation 3	Normal	Fail
				Force	Pass
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2 or Generation 3	Normal	Fail
				Force	Pass or fail1
	Auto max 4000	Auto	Generation 1 or Generation 2		
	Auto max 4000	Auto	Generation 3		

PortChannel Members	Configured Speed		New Member Type	Addition Type	Result
	PortChannel	New Member			
Generation 2 interfaces	Auto	Auto	Generation 1	Normal or force	Fail
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2 or Generation 3	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 2 or Generation 3	Normal	Fail
				Force	Pass
Generation 3 interfaces	Auto	Auto	Generation 1	Normal or force	Fail
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 2	Normal	Fail
				Force	Pass
	Auto max 2000	Auto	Generation 3	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 3	Normal	Fail
				Force	Pass
Generation 4 interfaces	Auto	Auto	Generation 1	Normal or force	Fail
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 2	Normal	Fail
				Force	Pass
	Auto max 2000	Auto	Generation 3 or Generation 4	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 3 or Generation 4	Normal	Fail
				Force	Pass

²⁰ If resources are not available.

Use the **show port-channel compatibility parameters** command to obtain information about PortChannel addition errors.

Default Settings

Table 37: Default Generation 2 Interface Parameters , on page 336 lists the default settings for Generation 2 interface parameters.

Table 37: Default Generation 2 Interface Parameters

Parameter	Default			
	48-Port 4-Gbps Switching Module	24-Port 4-Gbps Switching Module	12-Port 4-Gbps Switching Module	4-Port 10-Gbps Switching Module
Speed mode	auto ²¹	auto	auto	auto ²²
Rate mode	shared	shared	dedicated	dedicated
Port mode	Fx	Fx	auto ²³	auto ²⁴
BB_credit buffers	16	16	250	250
Performance buffers	—	—	145 ²⁵	1455

²¹ Auto speed mode on the 4-Gbps switching modules enables autosensing and negotiates to a maximum speed of 4 Gbps.

²² The 4-port 10-Gbps switching module only supports 10-Gbps traffic.

²³ Auto port mode on the 12-port 4-Gbps switching module interfaces can operate in E port mode, TE port mode, and Fx port mode.

²⁴ Auto port mode on the 4-port 10-Gbps switching module interfaces can operate in E port mode, TE port mode, and F port mode.

²⁵ Performance buffers are shared among all ports on the module.

Table 38: Default Generation 3 Interface Parameters, on page 336 lists the default settings for Generation 3 interface parameters.

Table 38: Default Generation 3 Interface Parameters

Parameter	Default		
	48-Port 8-Gbps Switching Module	24-Port 8-Gbps Switching Module	4/44-Port 8-Gbps Host-Optimized Switching Module
Speed mode ²⁶	auto	auto	auto_max_4G ²⁷
Rate mode	shared	shared	shared
Port mode	Fx	Fx	Fx
BB_credit buffers	32	32	32

²⁶ Auto speed mode on the 8-Gbps switching modules enables autosensing and negotiates to a maximum speed of 8 Gbps.

²⁷ Auto_max_4G speed mode on the 4/44-port 8-Gbps switching module negotiates to a maximum speed of 4 Gbps.

Table 39: Default Generation 4 Interface Parameters , on page 337 lists the default settings for Generation 4 interface parameters.

Table 39: Default Generation 4 Interface Parameters

Parameter	Default	
	48-Port 8-Gbps Advanced Fibre Channel Switching Module	32-Port 8-Gbps Advanced Fibre Channel Switching Module
Speed mode	auto ²⁸	auto
Rate mode	shared	shared
Port mode	Fx	Fx
BB_credit buffers	32	32

²⁸ Auto speed mode on the 8-Gbps switching modules enables autosensing and negotiates to a maximum speed of 8 Gbps.

Configuring Fibre Channel Interfaces

This section includes the following topics:

Task Flow for Migrating Interfaces from Shared Mode to Dedicated Mode

The 48-Port, 24-Port, and 4/44-Port 8-Gbps Fibre Channel switching modules support the following features:

- 1-Gbps, 2-Gbps, 4-Gbps, and 8-Gbps speed traffic
- Shared and dedicated rate mode
- ISL and Fx port modes
- Extended BB_credits

The 48-port and 24-port 4-Gbps Fibre Channel switching modules support the following features:

- 1-Gbps, 2-Gbps, and 4-Gbps speed traffic
- Shared and dedicated rate mode
- ISL (E or TE) and Fx (F or FL) port modes
- Extended BB_credits



Note

If you change the port bandwidth reservation parameters on a 48-port or 24-port 4-Gbps module, the change affects only the changed port. No other ports in the port group are affected.

To configure the 4-Gbps and 8-Gbps Fibre Channel switching modules when starting with the default configuration or when migrating from shared rate mode to dedicated rate mode, follow these steps:

Procedure

- Step 1** Take unused interfaces out of service to release resources for other interfaces, if necessary.
See the [Taking Interfaces Out of Service, on page 345](#).
- Step 2** Configure the traffic speed to use (1 Gbps, 2 Gbps, 4 Gbps, 8 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps).
See the [Dynamic Bandwidth Management, on page 320](#).
- Step 3** Configure the rate mode (dedicated or shared).
See the [Configuring Rate Mode, on page 341](#).
- Step 4** Configure the port mode.
See the Configuring Interface Modes section.
- Note** ISL ports cannot operate in shared rate mode.
- Step 5** Configure the BB_credits and extended BB_credits, as necessary.
See the Configuring Buffer-to-Buffer Credits section and the Configuring Extended BB_credits section.
-

Task Flow for Migrating Interfaces from Dedicated Mode to Shared Mode

To configure the 4-Gbps and 8-Gbps Fibre Channel switching modules migrating from dedicated rate mode to shared rate mode, follow these steps:

Procedure

- Step 1** Take unused interfaces out of service to release resources for other interfaces, if necessary.
See the [Taking Interfaces Out of Service, on page 345](#).
- Step 2** Configure the BB_credits and extended BB_credits, as necessary.
See the *Configuring Buffer-to-Buffer Credits* section, and the *Extended BB_credits on Generation 1 Switching Modules* section.
- Step 3** Configure the port mode.
See the *Configuring Interface Modes* section.
- Note** ISL ports cannot operate in shared rate mode.
- Step 4** Configure the rate mode (dedicated or shared) to use.
See the [Configuring Rate Mode, on page 341](#).
- Step 5** Configure the traffic speed (1 Gbps, 2 Gbps, 4 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps) to use.

See the [Dynamic Bandwidth Management, on page 320](#).

Task Flow for Configuring 12-Port 4-Gbps Module Interfaces

The 12-port 4-Gbps switching modules support the following features:

- 1-Gbps, 2-Gbps, and 4-Gbps speed traffic
- Only dedicated rate mode
- ISL (E or TE) and Fx (F or FL) port modes
- Extended BB_credits
- Performance buffers

To configure 4-port 10-Gbps switching modules when starting with the default configuration, follow these steps:

Procedure

- Step 1** Configure the traffic speed (1 Gbps, 2 Gbps, 4 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps) to use.
See [Dynamic Bandwidth Management, on page 320](#).
- Step 2** Configure the port mode.
See the *Configuring Interface Modes* section.
- Step 3** Configure the BB_credits, performance buffers, and extended BB_credits, as necessary.
See the *Configuring Buffer-to-Buffer Credits* section, and the *Configuring Extended BB_credits* section.
-

Task Flow for Configuring 4-Port 10-Gbps Module Interfaces

The 4-port 10-Gbps switching modules support the following features:

- Only 10-Gbps speed traffic
- Only dedicated rate mode
- ISL (E or TE) and F port modes
- Extended BB_credits
- Performance buffers

To configure 4-port 10-Gbps switching modules when starting with the default configuration, follow these steps:

Procedure

- Step 1** Configure the port mode.
See the *Configuring Interface Modes* section.

- Step 2** Configure the BB_credits, performance buffers, and extended BB_credits, as necessary.
See the *Configuring Buffer-to-Buffer Credits* section, and the *Configuring Extended BB_credits* section.

Reserving Bandwidth Quickly for the 8-Gbps Module Interfaces

Detailed Steps

To quickly reserve bandwidth for all the ports in the port groups on the Generation 3 Fibre Channel modules using the Device Manager, follow these steps:

Procedure

- Step 1** In the Device Manager window, right-click the 8-Gbps Fibre Channel module.

Figure 22: Device Manager - 8 Gbps Module - Pop-Up Menu

- Step 2** From the popup menu, select **Bandwidth Reservation Config...**

- Step 3** In the Bandwidth Reservation Configuration dialog box that is displayed, choose a bandwidth reservation scheme. ([Figure 23: RateMode Configuration Dialog Box, on page 340](#)).

Figure 23: RateMode Configuration Dialog Box

[Table 40: RateMode Configuration Schemes , on page 340](#) describes the default RateMode configuration schemes available in the Bandwidth Reservation Configuration dialog box for the 8-Gbps modules.

Table 40: RateMode Configuration Schemes

Module	Available RateMode Config Macros
DS-X9248-96K9 48-Port 8-Gbps Fibre Channel module	<ul style="list-style-type: none"> Dedicated 4 Gbps on the first port of each group and the remaining ports 8 Gbps shared Dedicated 8 Gbps on the first port of each group and the remaining ports 8 Gbps shared Shared 8 Gbps on all ports (initial and default settings)
DS-X9224-96K9 24-Port 8-Gbps Fibre Channel module	<ul style="list-style-type: none"> Dedicated 8 Gbps on the first port of each group and the remaining ports 8G shared Shared Auto²⁹ on all ports (initial and default settings)
DS-X9248-48K9 4/44-Port 8-Gbps Host-Optimized Fibre Channel module	<ul style="list-style-type: none"> Dedicated 2 Gbps on the first port of each group and the remaining ports 4 Gbps shared Dedicated 8 Gbps on the first port of each group and the remaining ports 4 Gbps shared Shared Auto with Maximum of 4 Gbps on all ports (initial and default settings)

²⁹ Auto is 8 Gbps.

- Step 4** Click **Apply**.
-

Configuring Port Speed

To configure dedicated bandwidth on an interface using DCNM-SAN, follow these steps:

Procedure

- Step 1** From the Fabric pane, select a switch or select a group of switches (SAN, fabric, VSAN) from the Logical Domains pane.
- Step 2** Expand **Switches**, expand **FC Interfaces** and select **Physical** from the Physical Attributes pane.
You see the **Physical > General** tab in the Interfaces pane.
- Step 3** Scroll until you see the row containing the switch and port you want to configure.
- Step 4** Select auto, 1Gb, 4Gb, or autoMax2G from the Speed Admin column.
- Note** The Generation 3, 8-Gbps Fibre Channel switching modules support the following speed configurations: 1G, 2G, 4G, 8G, autoMax2G, autoMax4G and the auto speed configuration configures autosensing for the interface with 8 Gbps of bandwidth reserved.
- The auto parameter enables autosensing on the interface. The autoMax2G parameter enables autosensing on the interface with a maximum speed of 2 Gbps.
- Note** If you change the port bandwidth reservation parameters on a 48-port or 24-port 4-Gbps, or any 8-Gbps Fibre Channel switching module, the change affects only the changed port. No other ports in the port group are affected.
- Step 5** Click the **Apply Changes** icon.
-

Configuring Rate Mode

To configure the rate mode (dedicated or shared) on an interface on a 4-Gbps or 8-Gbps Fibre Channel switching module using DCNM-SAN, follow these steps:

Procedure

- Step 1** Select a switch from the Fabric pane, or select a group of switches (SAN, fabric, VSAN) from the Logical Domains pane.
- Step 2** Expand **Switches > FC Interfaces** and then select **Physical** from the Physical Attributes pane.
You see the **Physical > General** tab in the Interfaces pane.
- Step 3** Scroll until you see the row containing the switch and port you want to configure.
- Step 4** Select **dedicated** or **shared** from the Rate Mode column.
- Step 5** Click the **Apply Changes** icon.
-

Configuring Local Switching



Note We recommend that you shut down all of the ports on the module before you execute the local switching command. If local switching is enabled, then ports cannot be configured in dedicated mode. If there are dedicated ports and you enter the local switching command, a warning is displayed and the operation is prevented.

Configuring Local Switching Using DCNM-SAN

To enable or disable local switching module using DCNM-SAN, follow these steps:

Procedure

- Step 1** Choose **Switches > Hardware**.
- Step 2** Click the **Module Config** tab. You see the Module Config dialog box.
- Step 3** Select a module and from the LocalSwitchingMode drop-down list, select **enabled** or **disabled**. This step either enables or disables the local switching for the selected module.
- Step 4** Click **Apply** to save the changes.

Configuring Local Switching Using Device Manager

To enable or disable local switching using Device Manager, follow these steps:

Procedure

- Step 1** Right-click a module and select **Configure**.
You see the Module dialog box. Click the **Config** tab.
- Step 2** Click the **enabled** or **disabled** radio button to enable or disable local switching in the selected module.
- Step 3** Click **Apply** to save the changes.

Disabling Restrictions on Oversubscription Ratios Using DCNM-SAN

To disable restrictions on oversubscription ratios on multiple 48-port or 24-port 4-Gbps, or any 8-Gbps Fibre Channel switching modules using DCNM-SAN, follow these steps:

Procedure

- Step 1** Choose **Switches > Hardware**.
- Step 2** Click the **Module Config** tab. You see the Module Config dialog box.

- Step 3** From the RateModeOversubscriptionLimit drop-down list, select **disabled** for each module for which you want to disable restrictions on oversubscription ratios.
- Step 4** Click **Apply** to save the changes.
-

Disabling Restrictions on Oversubscription Ratios Using Device Manager

To disable restrictions on oversubscription ratios on a single 48-port or 24-port 4-Gbps, or any 8-Gbps Fibre Channel switching module using Device Manager, follow these steps:

Procedure

- Step 1** Right-click a module and select **Configure**.
You see the Module dialog box.
- Step 2** Click the disabled radio button to disable restrictions on oversubscription ratios.
- Step 3** Click **Apply** to save the changes.
-

Enabling Restrictions on Oversubscription Ratios Using DCNM-SAN

To enable restrictions on over subscription ratios on multiple 48-port or 24-port 4-Gbps, or any 8-Gbps Fibre Channel switching modules using DCNM-SAN, follow these steps:

Procedure

- Step 1** Choose **Switches > Hardware**.
- Step 2** Click the **Module Config** tab. You see the Module Config dialog box.
- Step 3** From the RateMode Oversubscription Limit drop-down list, select enabled for each module for which you want to enable restrictions on oversubscription ratios.
- Step 4** Click **Apply** to save the changes.
-

Enabling Restrictions on Oversubscription Ratios Using Device Manager

To enable restrictions on over subscription ratios on a single 48-port or 24-port 4-Gbps, or any 8-Gbps Fibre Channel switching module using Device Manager, follow these steps:

Procedure

- Step 1** Right-click a module and select **Configure**.
You see the Module dialog box.
- Step 2** Click the **enabled** radio button to enable restrictions on oversubscription ratios.

- Step 3** Click **Apply** to save the changes.
-

Enabling Bandwidth Fairness Using DCNM-SAN

To enable bandwidth fairness on multiple 48-port or 24-port 4-Gbps, or any 8-Gbps Fibre Channel switching modules using DCNM-SAN, follow these steps:

Procedure

- Step 1** Choose **Switches > Hardware**.
- Step 2** Click the **Module Config** tab. You see the Module Config dialog box.
- Step 3** From the BandwidthFairness Admin drop-down list, select enable for each module for which you want to enable bandwidth fairness.
- Step 4** Click **Apply** to save the changes.
-

Enabling Bandwidth Fairness Using Device Manager

To enable bandwidth fairness on a single 48-port or 24-port 4-Gbps Fibre Channel switching module using Device Manager, follow these steps:

Procedure

- Step 1** Right-click a module and select **Configure**.
You see the Module dialog box.
- Step 2** Click the **enable** radio button to enable bandwidth fairness.
- Step 3** Click **Apply** to save the changes.
-

Disabling Bandwidth Fairness Using DCNM-SAN

To disable bandwidth fairness on multiple 48-port or 24-port 4-Gbps, or 8-Gbps Fibre Channel switching modules using DCNM-SAN, follow these steps:

Procedure

- Step 1** Choose **Switches > Hardware**.
- Step 2** Click the **Module Config** tab. You see the Module Config dialog box.
- Step 3** From the BandwidthFairness Admin drop-down list, select **disable** for each module for which you want to disable bandwidth fairness.
- Step 4** Click **Apply** to save the changes.
-

Disabling Bandwidth Fairness Using Device Manager

To disable bandwidth fairness on a single 48-port or 24-port 4-Gbps, or 8-Gbps Fibre Channel switching module using Device Manager, follow these steps:

Procedure

- Step 1** Right-click a module and select **Configure**.
You see the **Module** dialog box.
 - Step 2** Click the **disable** radio button to disable bandwidth fairness.
 - Step 3** Click **Apply** to save the changes.
-

Taking Interfaces Out of Service

To take an interface out of service using DCNM-SAN, follow these steps:

Procedure

- Step 1** Select a switch from the **Fabric** pane, or select a group of switches (**SAN**, **fabric**, **VSAN**) from the **Logical Domains** pane.
 - Step 2** Expand **Switches**, and expand **FC Interfaces > Physical** in the **Physical Attributes** pane.
 - Step 3** Click **General** tab. You see the **General** tab information in the **Information** pane.
 - Step 4** Scroll down until you see the row containing the switch and port you want to configure.
 - Step 5** Scroll right (if necessary) until you see the **Status Service** column.
 - Step 6** Select **in** or **out** from the **Status Service** column.
 - Step 7** Click the **Apply Changes** icon.
-

Releasing Shared Resources in a Port Group

To release the shared resources for a port group using DCNM-SAN, follow these steps:

Procedure

- Step 1** Select a switch from the **Fabric** pane, or select a group of switches (**SAN**, **fabric**, **VSAN**) from the **Logical Domains** pane.
- Step 2** Expand **Switches**, and expand **FC Interfaces > Physical** in the **Physical Attributes** pane.
- Step 3** Click **General** tab. You see the **General** tab information in the **Information** pane.
- Step 4** Scroll down until you see the row containing the switch and port you want to configure.
- Step 5** Scroll right (if necessary) until you see the **Status Service** column.
- Step 6** Select the status **out** from the **Status Service** column.

- Step 7** Click the **Apply Changes** icon.
- Step 8** Select the status **in** from the **Status Service** column.
- Step 9** Click the **Apply Changes** icon.

Verifying Fibre Channel Interfaces Configuration

To display Fibre Channel interface configuration information, perform one of the following tasks:

Command	Purpose
show module	Displays the module.
show module slot recovery-steps	Displays the slot for the module.
show port-resources module slot	Displays the port resources for the slot.
show interface fc slot/port	Displays the slot or port information.
show interface brief	Displays the interface.
show port index-allocation	Displays the port in the index allocation.
show port index-allocation startup	Displays the startup port in the index allocation.
show port-channel compatibility parameters	Displays the PortChannel compatibility parameters.
show module slot bandwidth-fairness	Displays the module slot bandwidth fairness information.

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS NX-OS Command Reference*.

Displaying Diagnostics for Multiple Ports

To view diagnostic information for multiple ports using Device Manager, follow these steps:

Procedure

- Step 1** Choose **Interface > FC All** and click the **Diagnostics** tab or hold down the **Control** key, and then click each port for which you want to view diagnostic information.
- Step 2** Right-click the selected ports, and then select **Configure**.
You see the FC Interfaces dialog box.
- Step 3** Click **Refresh** to view the latest diagnostic information.
- To view diagnostic information for a single port using Device Manager, follow these steps:
- Right-click a port, and then select **Configure**.
You see the port licensing options for the selected port.

- b) Click **Refresh** to view the latest information.
-



CHAPTER 13

Using the CFS Infrastructure

- [Monitoring Network Traffic Using SPAN, on page 349](#)

Monitoring Network Traffic Using SPAN

This chapter describes the Switched Port Analyzer (SPAN) features provided in switches in the Cisco MDS 9000 Family.

This chapter includes the following sections:

Information About SPAN

The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic through a Fibre Channel interface. Traffic through any Fibre Channel interface can be replicated to a special port called the SPAN destination port (SD port). Any Fibre Channel port in a switch can be configured as an SD port. Once an interface is in SD port mode, it cannot be used for normal data traffic. You can attach a Fibre Channel Analyzer to the SD port to monitor SPAN traffic.

SD ports do not receive frames, they only transmit a copy of the SPAN source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports (see [Figure 24: SPAN Transmission , on page 349](#)).

Figure 24: SPAN Transmission



This section covers the following topics:

SPAN Sources

SPAN sources refer to the interfaces from which traffic can be monitored. You can also specify VSAN as a SPAN source, in which case, all supported interfaces in the specified VSAN are included as SPAN sources. When a VSAN as a source is specified, then all physical ports and PortChannels in that VSAN are included as SPAN sources. You can choose the SPAN traffic in the ingress direction, the egress direction, or both directions for any source interface:

- Ingress source (Rx)—Traffic entering the switch fabric through this source interface is *spanned* or copied to the SD port (see [Figure 25: SPAN Traffic from the Ingress Direction , on page 350](#)).

Figure 25: SPAN Traffic from the Ingress Direction

- Egress source (Tx)—Traffic exiting the switch fabric through this source interface is spanned or copied to the SD port (see [Figure 26: SPAN Traffic from Egress Direction, on page 350](#)).

Figure 26: SPAN Traffic from Egress Direction

IPS Source Ports

SPAN capabilities are available on the IP Storage Services (IPS) module. The SPAN feature is only implemented on the FCIP and iSCSI virtual Fibre Channel port interfaces, not the physical Gigabit Ethernet ports. You can configure SPAN for ingress traffic, egress traffic, or traffic in both directions for all eight iSCSI and 24 FCIP interfaces that are available in the IPS module.



Note

You can configure SPAN for Ethernet traffic using Cisco switches or routers connected to the Cisco MDS 9000 Family IPS modules.

Allowed Source Interface Types

The SPAN feature is available for the following interface types:

- Physical ports such as F ports, FL ports, TE ports, E ports, and TL ports.
- Interface sup-fc0 (traffic to and from the supervisor):
 - The Fibre Channel traffic from the supervisor module to the switch fabric through the sup-fc0 interface is called ingress traffic. It is spanned when sup-fc0 is chosen as an ingress source port.
 - The Fibre Channel traffic from the switch fabric to the supervisor module through the sup-fc0 interface is called egress traffic. It is spanned when sup-fc0 is chosen as an egress source port.
- PortChannels
 - All ports in the PortChannel are included and spanned as sources.
 - You cannot specify individual ports in a PortChannel as SPAN sources. Previously configured SPAN-specific interface information is discarded.
- IPS module specific Fibre Channel interfaces:
 - iSCSI interfaces
 - FCIP interfaces

VSAN as a Source

SPAN sources refer to the interfaces from which traffic can be monitored. When a VSAN as a source is specified, then all physical ports and PortChannels in that VSAN are included as SPAN sources. A TE port is included only when the port VSAN of the TE port matches the source VSAN. A TE port is excluded even if the configured allowed VSAN list may have the source VSAN, but the port VSAN is different.

You cannot configure source interfaces (physical interfaces, PortChannels, or sup-fc interfaces) and source VSANs in the same SPAN session.

SPAN Sessions

Each SPAN session represents an association of one destination with a set of source(s) along with various other parameters that you specify to monitor the network traffic. One destination can be used by one or more SPAN sessions. You can configure up to 16 SPAN sessions in a switch. Each session can have several source ports and one destination port.

To activate any SPAN session, at least one source and the SD port must be up and functioning. Otherwise, traffic is not directed to the SD port.

**Tip**

A source can be shared by two sessions, however, each session must be in a different direction—one ingress and one egress.

You can temporarily deactivate (suspend) any SPAN session. The traffic monitoring is stopped during this time.

Specifying Filters

You can perform VSAN-based filtering to selectively monitor network traffic on specified VSANs. You can apply this VSAN filter to all sources in a session. Only VSANs present in the filter are spanned.

You can specify session VSAN filters that are applied to all sources in the specified session. These filters are bidirectional and apply to all sources configured in the session. Each SPAN session represents an association of one destination with a set of source(s) along with various other parameters that you specify to monitor the network traffic.

SD Port Characteristics

An SD port has the following characteristics:

- Ignores BB_credits.
- Allows data traffic only in the egress (Tx) direction.
- Does not require a device or an analyzer to be physically connected.
- Supports only 1 Gbps or 2 Gbps speeds. The auto speed option is not allowed.
- Multiple sessions can share the same destination ports.
- If the SD port is shut down, all shared sessions stop generating SPAN traffic.
- The outgoing frames can be encapsulated in Extended Inter-Switch Link (EISL) format.
- The SD port does not have a port VSAN.
- SD ports cannot be configured using Storage Services Modules (SSMs).
- The port mode cannot be changed if it is being used for a SPAN session.

**Note**

If you need to change an SD port mode to another port mode, first remove the SD port from all sessions and then change the port mode.

Monitoring Traffic Using Fibre Channel Analyzers

You can use SPAN to monitor traffic on an interface without any traffic disruption. This feature is especially useful in troubleshooting scenarios in which traffic disruption changes the problem environment and makes it difficult to reproduce the problem. You can monitor traffic in either of the following two ways:

- Without SPAN
- With SPAN

Monitoring Without SPAN

You can monitor traffic using interface fc1/1 in a Cisco MDS 9000 Family switch that is connected to another switch or host. You need to physically connect a Fibre Channel analyzer between the switch and the storage device to analyze the traffic through interface fc1/1 (see [Figure 27: Fibre Channel Analyzer Usage Without SPAN](#), on page 352).

Figure 27: Fibre Channel Analyzer Usage Without SPAN



This type of connection has the following limitations:

- It requires you to physically insert the FC analyzer between the two network devices.
- It disrupts traffic when the Fibre Channel analyzer is physically connected.
- The analyzer captures data only on the Rx links in both port 1 and port 2. Port 1 captures traffic exiting interface fc1/1 and port 2 captures ingress traffic into interface fc1/1.

Monitoring with SPAN

Using SPAN you can capture the same traffic scenario without any traffic disruption. The Fibre Channel analyzer uses the ingress (Rx) link at port 1 to capture all the frames going out of the interface fc1/1. It uses the ingress link at port 2 to capture all the ingress traffic on interface fc1/1.

Using SPAN you can monitor ingress traffic on fc1/1 at SD port fc2/2 and egress traffic on SD port fc2/1. This traffic is seamlessly captured by the FC analyzer (see [Figure 28: Fibre Channel Analyzer Using SPAN](#), on page 352).

Figure 28: Fibre Channel Analyzer Using SPAN



Single SD Port to Monitor Traffic

You do not need to use two SD ports to monitor bidirectional traffic on any interface. You can use one SD port and one FC analyzer port by monitoring traffic on the interface at the same SD port fc2/1.

[Figure 29: Fibre Channel Analyzer Using a Single SD Port](#), on page 352 shows a SPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress directions. This setup is more advantageous and cost effective than the setup shown in the *Monitoring with SPAN* section. It uses one SD port and one port on the analyzer, instead of using a full, two-port analyzer.

Figure 29: Fibre Channel Analyzer Using a Single SD Port



To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

SD Port Configuration

The SD port in the destination switch enables the FC analyzer to receive the RSPAN traffic from the Fibre Channel tunnel. [Figure 30: RSPAN Tunnel Configuration, on page 353](#) depicts an RSPAN tunnel configuration, now that tunnel destination is also configured.

Figure 30: RSPAN Tunnel Configuration



Note SD ports cannot be configured using Storage Services Modules (SSMs).

Mapping the FC Tunnel

The **tunnel-id-map** option specifies the egress interface of the tunnel at the destination switch (see [Figure 55-8](#)).



Creating VSAN Interfaces

[Figure 31: FC Tunnel Configuration, on page 353](#) depicts a basic FC tunnel configuration.

Figure 31: FC Tunnel Configuration



Note This example assumes that VSAN 5 is already configured in the VSAN database.

Remote SPAN



Note Remote SPAN is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeSystem.

The Remote SPAN (RSPAN) feature enables you to remotely monitor traffic for one or more SPAN sources distributed in one or more source switches in a Fibre Channel fabric. The SPAN destination (SD) port is used for remote monitoring in a destination switch. A destination switch is usually different from the source switch(es) but is attached to the same Fibre Channel fabric. You can replicate and monitor traffic in any remote Cisco MDS 9000 Family switch or director, just as you would monitor traffic in a Cisco MDS source switch.

The RSPAN feature is nonintrusive and does not affect network traffic switching for those SPAN source ports. Traffic captured on the remote switch is tunneled across a Fibre Channel fabric which has trunking enabled

on all switches in the path from the source switch to the destination switch. The Fibre Channel tunnel is structured using trunked ISL (TE) ports. In addition to TE ports, the RSPAN feature uses two other interface types (see [Figure 32: RSPAN Transmission](#), on page 354):

- SD ports—A passive port from which remote SPAN traffic can be obtained by the FC analyzer.
- ST ports—A SPAN tunnel (ST) port is an entry point port in the source switch for the RSPAN Fibre Channel tunnel. ST ports are special RSPAN ports and cannot be used for normal Fibre Channel traffic.

Figure 32: RSPAN Transmission



Advantages of Using RSPAN

The RSPAN features has the following advantages:

- Enables nondisruptive traffic monitoring at a remote location.
- Provides a cost effective solution by using one SD port to monitor remote traffic on multiple switches.
- Works with any Fibre Channel analyzer.
- Is compatible with the Cisco MDS 9000 Port Analyzer adapters.
- Does not affect traffic in the source switch, but shares the ISL bandwidth with other ports in the fabric.

FC and RSPAN Tunnels

An FC tunnel is a logical data path between a source switch and a destination switch. The FC tunnel originates from the source switch and terminates at the remotely located destination switch.

RSPAN uses a special Fibre Channel tunnel (FC tunnel) that originates at the ST port in the source switch and terminates at the SD port in the destination switch. You must bind the FC tunnel to an ST port in the source switch and map the same FC tunnel to an SD port in the destination switch. Once the mapping and binding is configured, the FC tunnel is referred to as an RSPAN tunnel (see [Figure 33: FC and RSPAN Tunnel](#), on page 354).

Figure 33: FC and RSPAN Tunnel



ST Port Configuration

Once the FC tunnel is created, be sure to configure the ST port to bind it to the FC tunnel at the source switch. The FC tunnel becomes an RSPAN tunnel once the binding and mapping is complete.

[Figure 34: Binding the FC Tunnel](#), on page 354 depicts a basic FC tunnel configuration.

Figure 34: Binding the FC Tunnel



ST Port Characteristics

ST ports have the following characteristics:

- ST ports perform the RSPAN encapsulation of the FC frame.
- ST ports do not use BB_credits.

- One ST port can only be bound to one FC tunnel.
- ST ports cannot be used for any purpose other than to carry RSPAN traffic.
- ST ports cannot be configured using Storage Services Modules (SSMs).

Creating Explicit Paths

You can specify an explicit path through the Cisco MDS Fibre Channel fabric (source-based routing), using the **explicit-path** option. For example, if you have multiple paths to a tunnel destination, you can use this option to specify the FC tunnel to always take one path to the destination switch. The software then uses this specified path even if other paths are available.

This option is especially useful if you prefer to direct the traffic through a certain path although other paths are available. In an RSPAN situation, you can specify the explicit path so the RSPAN traffic does not interfere with the existing user traffic. You can create any number of explicit paths in a switch (see [Figure 35: Explicit Path Configuration, on page 355](#)).

Figure 35: Explicit Path Configuration



Guidelines and Limitations

SPAN Configuration Guidelines

The following guidelines and limitations apply for SPAN configurations:

- You can configure up to 16 SPAN sessions with multiple ingress (Rx) sources.
- You can configure a maximum of three SPAN sessions with one egress (Tx) port.
- In a 32-port switching module, you must configure the same session in all four ports in one port group (unit). If you wish, you can also configure only two or three ports in this unit.
- SPAN frames are dropped if the sum of the bandwidth of the sources exceeds the speed of the destination port.
- Frames dropped by a source port are not spanned.
- SPAN does not capture pause frames in a Fibre Channel over Ethernet (FCoE) network because pause frames sent from the virtual expansion (VE) port are generated and terminated by the outermost MAC layer. For more information on FCoE, see the Cisco NX-OS FCoE Configuration Guide for Cisco Nexus 7000 and Cisco MDS 9500.

Guidelines to Configure VSANs as a Source

The following guidelines apply when configuring VSANs as a source:

- Traffic on all interfaces included in a source VSAN is spanned only in the ingress direction.
- If a VSAN is specified as a source, you cannot perform interface-level SPAN configuration on the interfaces that are included in the VSAN. Previously configured SPAN-specific interface information is discarded.
- If an interface in a VSAN is configured as a source, you cannot configure that VSAN as a source. You must first remove the existing SPAN configurations on such interfaces before configuring VSAN as a source.
- Interfaces are only included as sources when the port VSAN matches the source VSAN. [Figure 36: VSAN as a Source , on page 356](#) displays a configuration using VSAN 2 as a source:

- All ports in the switch are in VSAN 1 except fc1/1.
- Interface fc1/1 is the TE port with port VSAN 2. VSANs 1, 2, and 3 are configured in the allowed list.
- VSAN 1 and VSAN 2 are configured as SPAN sources.

Figure 36: VSAN as a Source



For this configuration, the following apply:

- VSAN 2 as a source includes only the TE port fc1/1 that has port VSAN 2.
- VSAN 1 as a source does not include the TE port fc1/1 because the port VSAN does not match VSAN 1.

Guidelines to Specifying Filters

The following guidelines apply to SPAN filters:

- PortChannel configurations are applied to all ports in the PortChannel.
- If no filters are specified, the traffic from all active VSANs for that interface is spanned by default.
- While you can specify arbitrary VSAN filters in a session, traffic can only be monitored on the port VSAN or on allowed-active VSANs in that interface.

RSPAN Configuration Guidelines

The following guidelines apply for a SPAN configuration:

- All switches in the end-to-end path of the RSPAN tunnel must belong to the Cisco MDS 9000 Family.
- All VSANs with RSPAN traffic must be enabled. If a VSAN containing RSPAN traffic is not enabled, it is dropped.
- The following configurations must be performed on *each* switch in the end-to-end path of the Fibre Channel tunnel in which RSPAN is to be implemented:
 - Trunking must be enabled (enabled by default) and the trunk enabled link must be the lowest cost link in the path.
 - VSAN interface must be configured.
 - The Fibre Channel tunnel feature must be enabled (disabled by default).
 - IP routing must be enabled (disabled by default).



Note

If the IP address is in the same subnet as the VSAN, the VSAN interface does not have to be configured for all VSANs on which the traffic is spanned.

- A single Fibre Channel switch port must be dedicated for the ST port functionality.
- Do not configure the port to be monitored as the ST port.
- The FC tunnel's IP address must reside in the same subnet as the VSAN interface.

Default SPAN and RSPAN Settings

[Table 41: Default SPAN Configuration Parameters](#), on page 357 lists the default settings for SPAN parameters.

Table 41: Default SPAN Configuration Parameters

Parameters	Default
SPAN session	Active.
If filters are not specified	SPAN traffic includes traffic through a specific interface from all active VSANs.
Encapsulation	Disabled.
SD port	Output frame format is Fibre Channel.

[Table 42: Default RSPAN Configuration Parameters](#), on page 357 lists the default settings for RSPAN parameters.

Table 42: Default RSPAN Configuration Parameters

Parameters	Default
FC tunnel	Disabled.
Explicit path	Not configured.
Minimum cost path	Used if explicit path is not configured.

Configuring SPAN

The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic through a Fibre Channel interface.

This section covers the following topics:

Configuring SD Ports for SPAN

To monitor network traffic using SD ports, follow these steps:

Procedure

-
- | | |
|---------------|---------------------------------------------------------------------|
| 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. |
-

Configuring SD Ports for SPAN using DM

To configure an SD port for SPAN monitoring using Device Manager, follow these steps:

Procedure

- Step 1** Right-click the port you want to configure and select **Configure**.
You see the general port configuration dialog.
- Step 2** Under Mode, choose **SD**.
- Step 3** Click **Apply** to accept the change.
- Step 4** Close the dialog box.
-

Configuring SPAN max-queued-packets

When a SPAN destination port is oversubscribed or has more source traffic than the speed of the destination port, the source ports of the SPAN session will reduce in their throughput. The impact is proportional to the amount of source traffic flowing in. Lowering the max-queued-packets value from the default value of 15 to 1 prevents the impact on the source ports. It is necessary to reconsider the default value for this setting as it may impact the source interface throughput.

The following are the requirements:

- The span max-queued-packets can be changed only if no SPAN sessions are currently active on the switch.
- If you are spanning the traffic going through an FCIP interface, SPAN copies may be dropped even if the SD interface has more bandwidth than the amount of traffic being replicated. To avoid SPAN drops, set the max-queued-packets to a higher value; for example, 100.

By default, SPAN frames are dropped if the sum of the bandwidth of the source interfaces exceed the bandwidth of the destination port. With a higher value, the SPAN traffic has a higher probability of reaching the SPAN destination port instead of being dropped at the expense of data traffic throughput.

Creating SPAN Sessions

To create SPAN sessions, follow these steps:

Procedure

- 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.
- Step 4** Choose the session ID (the ID range may vary depending on platform type and version) using the up or down arrows and click Create.
- Step 5** Repeat Step 4 for each session 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.

Configuring SPAN for Generation 2 Fabric Switches

Cisco Generation 2 fabric switches (such as MDS 9124) support SPAN sessions in both directions, Rx and Tx.



Note While using Generation 2 fabric switches, you cannot create an additional active SPAN session when you already have one.

the following are the restrictions:

- You can specify multiple SPAN source interfaces in Rx and Tx directions. However, the direction should be explicitly mentioned at the end of the command. The SPAN will reject any source interface configuration that fails to mention the direction.
- You cannot mix ingress and egress interfaces in the same SPAN session. The SPAN will reject any configuration that mixes Rx and Tx directions. However, you can specify multiple SPAN source interfaces in a single direction.

Editing SPAN Sources

To edit a SPAN source, follow these steps:

Procedure

- Step 1** Choose Interface > SPAN.
You see the SPAN dialog box.
- Step 2** Click the 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.
You see the 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** Click the direction (**receive** or **transmit**) 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

To delete a SPAN session, follow these steps:

Procedure

- Step 1** Choose Interface > SPAN.
You see the SPAN dialog box.
- Step 2** Click the Sessions tab.
- Step 3** Click the SPAN session you want to delete.
- Step 4** Click Delete.
The SPAN session is deleted.
- Step 5** Close the dialog box.
-

Encapsulating Frames

The frame encapsulation feature is disabled by default. If you enable the encapsulation feature, all outgoing frames are encapsulated.

Configuring Fibre Channel Analyzers Using SPAN

To configure Fibre Channel Analyzers using SPAN for the example in the *Fibre Channel Analyzer Using SPAN* section, follow these steps:

Procedure

- Step 1** Configure SPAN on interface fc1/1 in the ingress (Rx) direction to send traffic on SD port fc2/1 using session 1.
- Step 2** Configure SPAN on interface fc1/1 in the egress (Tx) direction to send traffic on SD port fc2/2 using session 2.
- Step 3** Physically connect fc2/1 to port 1 on the Fibre Channel analyzer.
- Step 4** Physically connect fc2/2 to port 2 on the Fibre Channel analyzer.
-

Configuring RSPAN

The RSPAN tunnel begins in the source switch and terminates in the destination switch. This section assumes Switch S to be the source and Switch D to be the destination.

The following are the prerequisites:

- In addition to the source and destination switches, the VSAN must also be configured in each Cisco MDS switch in the Fibre Channel fabric, if they exist.

To monitor network traffic using the RSPAN feature, follow these steps:

Procedure

-
- | | |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Create VSAN interfaces in destination switch (Switch D) and source switch (Switch S) to facilitate the Fibre Channel tunnel (FC tunnel) creation. |
| Step 2 | Enable the FC tunnel in each switch in the end-to-end path of the tunnel. |
| Step 3 | Initiate the FC tunnel (in Switch S) and map the tunnel to the VSAN interface's IP address (in Switch D) so all RSPAN traffic from the tunnel is directed to the SD port. |
| Step 4 | Configure SD ports for SPAN monitoring in the destination switch (Switch D). |
| Step 5 | Configure the ST port in the source switch (Switch S) and bind the ST port to the FC tunnel. |
| Step 6 | Create an RSPAN session in the source switch (in Switch S) to monitor network traffic. |
-

Configuring the Source Switch

This section identifies the tasks that must be performed in the source switch (Switch S):

Enabling FC Tunnels

The following are the restrictions:

- FC tunnels do not work over nontrunking ISLs.
- The interface cannot be operationally up until the FC tunnel mapping is configured in the destination switch.



Note Be sure to enable this feature in each switch in the end-to-end path in the fabric.

Configuring All Intermediate Switches

This section identifies the tasks that must be performed in all intermediate switches in the end-to-end path of the RSPAN tunnel:

Configuring VSAN Interfaces

An RSPAN tunnel configuration is terminated in the destination switch (Switch D).



Note This example assumes that VSAN 5 is already configured in the VSAN database.

Enabling IP Routing

The IP routing feature is disabled by default. Be sure to enable IP routing in each switch (including the source and destination switches) in the end-to-end path in the fabric. This procedure is required to set up the FC tunnel.

Configuring the Destination Switch

This section identifies the tasks that must be performed in the destination switch (Switch D):

Configuring the SD Port

SD ports cannot be configured using Storage Services Modules (SSMs).

Monitoring RSPAN Traffic

Once the session is configured, other SPAN sources for this session can also be configured as required. [Figure 37: Fibre Channel Analyzer Using a Single SD Port to Monitor RSPAN Traffic, on page 362](#) shows an RSPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress directions.

Figure 37: Fibre Channel Analyzer Using a Single SD Port to Monitor RSPAN Traffic



To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

Configuration Examples for RSPAN

This section covers the following topic:



Note

RSPAN can be combined with the local SPAN feature so SD ports forward local SPAN traffic along with remote SPAN traffic. Various SPAN source and tunnel scenarios are described in this section.

Single Source with One RSPAN Tunnel

The source Switch S and the destination Switch D are interconnected through a Fibre Channel fabric. An RSPAN tunnel is configured as a destination interface for the SPAN session and the ST port forwards SPAN traffic through the RSPAN tunnel (see [Figure 38: RSPAN Scenario with One Source Switch, One Destination Switch, and One Tunnel, on page 362](#)).

Figure 38: RSPAN Scenario with One Source Switch, One Destination Switch, and One Tunnel



Single Source with Multiple RSPAN Tunnels

[Figure 39: RSPAN Scenario with One Source Switch, One Destination Switch, and Multiple Tunnels, on page 363](#) displays two separate RSPAN tunnels configured between Switches S and N. Each tunnel has an associated ST port in the source switch and a separate SD port in the destination switch. This configuration is useful for troubleshooting purposes.

Figure 39: RSPAN Scenario with One Source Switch, One Destination Switch, and Multiple Tunnels



Multiple Sources with Multiple RSPAN Tunnels

Figure 40: RSPAN Scenario with Two Source Switches, a Destination Switch, and Multiple Tunnels, on page 363 displays two separate RSPAN tunnels configured between Switches S1 and S2. Both tunnels have an associated ST port in their respective source switch and terminate in the same SD port in the destination switch.

Figure 40: RSPAN Scenario with Two Source Switches, a Destination Switch, and Multiple Tunnels



This configuration is useful for remote monitoring purposes. For example, the administrator may be at the destination switch and can remotely monitor the two source switches.

Field Descriptions for SPAN

This section describes the field descriptions for SPAN.

SPAN Sessions

Field	Description
Dest Interface	The Span Destination port interface.
Filter VSAN List	The VSANs that are assigned to this session.
Status Admin	Suspend an active session or activate an inactive session.
Status Oper	The current state of the session.
Description	The description of the session status.
VSAN List	The VSANs that are assigned to this session.
Or Interface (Direction)	The destination port ID to be configured for the session.
Inactive Reason	Description of the reason why this session is not active.

Related Topics

[SPAN Sessions, on page 351](#)

[Creating SPAN Sessions, on page 358](#)

[Deleting SPAN Sessions, on page 360](#)

[Information About SPAN, on page 349](#)

[Editing SPAN Sources, on page 359](#)

Span Global

Field	Description
MaxQueuedSpanPackets	This field specifies the drop threshold packets for all span sessions. The MaxQueuedSpanPackets field is only available when no session is active.

SPAN Source Interfaces

Field	Description
Interface, Direction	The destination port ID configured for the session, and the direction of traffic.



CHAPTER 14

Configuring SNMP

- [Configuring SNMP, on page 365](#)

Configuring SNMP

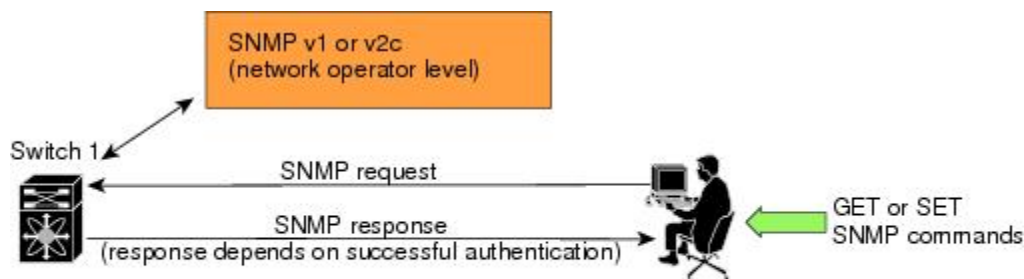
The CLI and SNMP use common roles in all switches in the Cisco MDS 9000 Family. You can use SNMP to modify a role that was created using the CLI and vice versa.

Users, passwords, and roles for all CLI and SNMP users are the same. A user configured through the CLI can access the switch using SNMP (for example, the DCNM-SAN or the Device Manager) and vice versa.

Information About SNMP Security

SNMP is an application layer protocol that facilitates the exchange of management information between network devices. In all Cisco MDS 9000 Family switches, three SNMP versions are available: SNMPv1, SNMPv2c, and SNMPv3.

Figure 41: SNMP Security



41-8

This section includes the following topics:

SNMP Version 1 and Version 2c

SNMP Version 1 (SNMPv1) and SNMP Version 2c (SNMPv2c) use a community string match for user authentication. Community strings provided a weak form of access control in earlier versions of SNMP.

SNMPv3 provides much improved access control using strong authentication and should be preferred over SNMPv1 and SNMPv2c wherever it is supported.

SNMP Version 3

SNMP Version 3 (SNMPv3) is an interoperable standards-based protocol for network management. SNMPv3 provides secure access to devices by a combination of authenticating and encrypting frames over the network. The security features provided in SNMPv3 are:

- Message integrity—Ensures that a packet has not been tampered with in-transit.
- Authentication—Determines the message is from a valid source.
- Encryption—Scrambles the packet contents to prevent it from being seen by unauthorized sources.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

SNMPv3 CLI User Management and AAA Integration

The Cisco NX-OS software implements RFC 3414 and RFC 3415, including user-based security model (USM) and role-based access control. While SNMP and the CLI have common role management and share the same credentials and access privileges, the local user database was not synchronized in earlier releases.

SNMPv3 user management can be centralized at the AAA server level. This centralized user management allows the SNMP agent running on the Cisco MDS switch to leverage the user authentication service of the AAA server. Once user authentication is verified, the SNMP PDUs are processed further. The AAA server also is used to store user group names. SNMP uses the group names to apply the access/role policy that is locally available in the switch.

Restricting Switch Access

You can restrict access to a Cisco MDS 9000 Family switch using IP access control lists (IP-ACLs).

Group-Based SNMP Access



Note Because *group* is a standard SNMP term used industry-wide, we refer to role(s) as group(s) in this SNMP section.

SNMP access rights are organized by groups. Each group in SNMP is similar to a role through the CLI. Each group is defined with three accesses: read access, write access, and notification access. Each access can be enabled or disabled within each group.

You can begin communicating with the agent once your user name is created, your roles are set up by your administrator, and you are added to the roles.

Creating and Modifying Users

You can create users or modify existing users using SNMP, DCNM-SAN, or the CLI.

- SNMP—Create a user as a clone of an existing user in the `usmUserTable` on the switch. Once you have created the user, change the cloned secret key before activating the user. Refer to RFC 2574.
- DCNM-SAN.
- CLI—Create a user or modify an existing user using the **`snmp-server user`** command.

A network-operator and network-admin roles are available in a Cisco MDS 9000 Family switch. There is also a default-role if you want to use the GUI (DCNM-SAN and Device Manager). You can also use any role that is configured in the Common Roles database.



Tip All updates to the CLI security database and the SNMP user database are synchronized. You can use the SNMP password to log into either DCNM-SAN or Device Manager. However, after you use the CLI password to log into DCNM-SAN or Device Manager, you must use the CLI password for all future logins. If a user exists in both the SNMP database and the CLI database before upgrading to Cisco MDS SAN-OS Release 2.0(1b), then the set of roles assigned to the user becomes the union of both sets of roles after the upgrade.

AES Encryption-Based Privacy

The Advanced Encryption Standard (AES) is the symmetric cipher algorithm. The Cisco NX-OS software uses AES as one of the privacy protocols for SNMP message encryption and conforms with RFC 3826.

The **`priv`** option offers a choice of DES or 128-bit AES encryption for SNMP security encryption. The **`priv`** option along with the **`aes-128`** token indicates that this privacy password is for generating a 128-bit AES key. The AES **`priv`** password can have a minimum of eight characters. If the passphrases are specified in clear text, you can specify a maximum of 64 characters. If you use the localized key, you can specify a maximum of 130 characters.



Note For an SNMPv3 operation using the external AAA server, user configurations in the external AAA server require AES to be the privacy protocol to use SNMP PDU encryption.

Enabling SNMP Notifications

Notifications (traps and informs) are system alerts that the switch generates when certain events occur. You can enable or disable notifications. By default, no notification is defined or issued. If a notification name is not specified, all notifications are disabled or enabled.

With the SNMP central infra feature, you can add the traps that need to be enabled or disabled. The MIB CISCO-NOTIFICATION-CONTROL-MIB is supported to enable the use of a MIB browser to control notification generation.

LinkUp/LinkDown Notifications for Switches

You can configure which LinkUp/LinkDown notifications to enable on switches. You can enable the following types of LinkUp/LinkDown notifications:

- Cisco—Only notifications (`cieLinkUp`, `cieLinkDown`) defined in CISCO-IF-EXTENSION-MIB.my are sent for an interface, if `ifLinkUpDownTrapEnable` (defined in IF-MIB) is enabled for that interface.

- IETF—Only notifications (LinkUp, LinkDown) defined in IF-MIB are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. Only the varbinds defined in the notification definition are sent with the notifications.
- IEFT extended—Only notifications (LinkUp, LinkDown) defined in IF-MIB are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. In addition to the varbinds defined in the notification definition, varbinds defined in the IF-MIB specific to the Cisco Systems implementation are sent. This is the default setting.
- IEFT Cisco—Only notifications (LinkUp, LinkDown) defined in IF-MIB and notifications (cieLinkUp, cieLinkDown) defined in CISCO-IF-EXTENSION-MIB.my are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. Only the varbinds defined in the notification definition are sent with the linkUp and linkDown notifications.
- IEFT extended Cisco—Only notifications (LinkUp, LinkDown) defined in IF-MIB and notifications (cieLinkUp, cieLinkDown) defined in CISCO-IF-EXTENSION-MIB.my are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. In addition to the varbinds defined in linkUp and linkDown notification definition, varbinds defined in the IF-MIB specific to the Cisco Systems implementation are sent with the LinkUp and LinkDown notifications.

**Note**

For more information on the varbinds defined in the IF-MIB specific to the Cisco Systems implementation, refer to the [Cisco MDS 9000 Family MIB Quick Reference](#).

Scope of LinkUp and LinkDown Trap Settings

The LinkUp and LinkDown trap settings for the interfaces generate traps based on the following scope:

Switch-level Trap Setting	Interface-level Trap Setting	Trap Generated for Interface Links?
Enabled (default)	Enabled (default)	Yes
Enabled	Disabled	No
Disabled	Enabled	No
Disabled	Disabled	No

Default Settings

The following table lists the default settings for all SNMP features in any switch.

Table 43: Default SNMP Settings

Parameters	Default
User account	No expiry (unless configured)
Password	None

Configuring SNMP

SNMP is an application layer protocol that facilitates the exchange of management information between network devices.

Assigning SNMP Switch Contact and Location Information

You can assign the switch contact information, which is limited to 32 characters (without spaces), and the switch location.

To configure contact and location information, follow these steps:

Procedure

-
- | | |
|---------------|------------------------------------------------------------------------------------------------------------|
| Step 1 | Expand Switches from the Physical Attributes pane.
You see the switch settings in the Information pane. |
| Step 2 | Fill in the Location and Contact fields for each switch. |
| Step 3 | Click Apply Changes to save these changes or click Undo Changes to discard any unsaved changes. |
-

Configuring SNMP Users from the CLI

The passphrase specified in the **snmp-server user** command and the **username** command are synchronized.

Restrictions

- Avoid using the **localizedkey** option when configuring an SNMP user from the CLI. The localized keys are not portable across devices as they contain device engine ID information. If a configuration file is copied to the device, the passwords may not be set correctly if the configuration file was generated at a different device. Explicitly configure the desired passwords after copying the configuration into the device. Passwords specified with the **localizedkey** option are limited to 130 characters.



Note The **snmp-server user** command takes the engineID as an additional parameter. The engineID creates the notification target user (see the [Configuring the Notification Target User](#), on page 376). If the engineID is not specified, the local user is created.

Enforcing SNMPv3 Message Encryption

By default the SNMP agent allows the securityLevel parameters of authNoPriv and authPriv for the SNMPv3 messages that use user-configured SNMPv3 message encryption with auth and priv keys.



Note Either to create a new SNMPv3 user or modify password of SNMPv3 user, the DCNM login user need to have enabled with DES/AES privacy password. Since the creating and modifying SNMP SET request need to be encrypted, the login user password needs to have the privacy password.

To enforce the message encryption for a user, follow these steps:

Procedure

- Step 1** Expand Switches, expand Security, and then select Users and Roles from the Physical Attributes pane.
 - Step 2** Click the Users tab in the Information pane to see a list of users.
 - Step 3** Click Create Row.
You see the Create Users dialog box.
 - Step 4** Enter the user name in the New User field.
 - Step 5** Select the role from the Role drop-down menu. You can enter a new role name in the field if you do not want to select one from the drop-down menu. If you do this, you must go back and configure this role appropriately.
 - Step 6** Enter a password for the user in Password field.
 - Step 7** Click the **Privacy** tab.
 - Step 8** Check the Enforce SNMP Privacy Encryption check box to encrypt management traffic.
 - Step 9** Click Create to create the new entry.
-

Enforce the SNMPv3 message encryption globally

To enforce the SNMPv3 message encryption globally on all the users, follow these steps:

Procedure

- Step 1** Select a VSAN in the Logical Domains pane. This will not work if you select All VSANS.
 - Step 2** Expand Switches, expand Security, and then select Users and Roles in the Physical Attributes pane. Click the Global tab in the Information pane.
 - Step 3** Check the GlobalEnforcePriv check box.
 - Step 4** Click the Apply Changes icon to save these changes.
-

Assigning SNMPv3 Users to Multiple Roles

The SNMP server user configuration is enhanced to accommodate multiple roles (groups) for SNMPv3 users. After the initial SNMPv3 user creation, you can map additional roles for the user.

- Only users belonging to a network-admin role can assign roles to other users.

To add multiple roles to a new user, follow these steps:

Procedure

- Step 1** Expand Switches, expand Security, and then select Users and Roles from the Physical Attributes pane.
- Step 2** Click the Users tab in the Information pane to see a list of users.

- Step 3** Click Create Row.
You see the Create Users dialog box.
- Step 4** Choose roles using the check boxes.
- Step 5** Choose an option for Digest and one for Encryption.
- Step 6** (Optional) Provide an expiration date for the user and the file name of an SSH key.
- Step 7** Click Create to create the new roles.
-

Adding or Deleting Communities

You can configure read-only or read-write access for SNMPv1 and SNMPv2 users. Refer to RFC 2576.
To create an SNMPv1 or SNMPv2c community string, follow these steps:

Procedure

- Step 1** Expand Switches, expand Security, and then select Users and Roles from the Physical Attributes pane.
- Step 2** Click the Communities tab in the Information pane.
You see the existing communities.
- Step 3** Click Create Row.
You see the Create Community String dialog box.
- Step 4** Check the Switch check boxes to specify one or more switches.
- Step 5** Enter the community name in the Community field.
- Step 6** Select the role from Role drop-down list.
- Note** You can enter a new role name in the field if you do not want to select one from the drop-down list.
If you do this, you must go back and configure this role appropriately.
- Step 7** Click Create to create the new entry.
-

Deleting a Community String

To delete a community string, follow these steps:

Procedure

- Step 1** Expand Switches, expand Security, and then select Users and Roles from the Physical Attributes pane.
- Step 2** Click the Communities tab in the Information pane.
- Step 3** Click the name of the community you want to delete.
- Step 4** Click Delete Row to delete this community.
-

Configuring SNMP Trap and Inform Notifications

You can configure the Cisco MDS switch to send notifications to SNMP managers when particular events occur.



Note You must enable the RMON traps in the SNMP configuration. For more information, refer to *Configuring SNMP*.



Note Use the SNMP-TARGET-MIB to obtain more information on the destinations to which notifications are to be sent either as traps or as informs. Refer to the [Cisco MDS 9000 Family MIB Quick Reference](#).

Configuring SNMPv2c Notifications

To configure SNMPv2c notifications, follow these steps:

Procedure

- Step 1** Expand Events and then select SNMP Traps in the Physical Attributes pane.
You see the SNMP notification configuration in the Information pane.
- Step 2** Click the Destinations tab to add or modify a receiver for SNMP notifications.
- Step 3** Click Create Row to create a new notification destination.
You see the Create Destinations dialog box.
- Step 4** Check the switches for which you want to configure a new destination.
- Step 5** Set the destination IP address and UDP port.
- Step 6** Choose either the trap or inform radio button.
- Step 7** (Optional) Set the timeout or retry count values.
- Step 8** Click Create to add this destination to the selected switches.
- Step 9** (Optional) Click the Other tab to enable specific notification types per switch.
- Step 10** Click the Apply changes icon to create the entry.

What to do next



Note Switches can forward events (SNMP traps and informs) up to 10 destinations.

Configuring SNMPv3 Notifications

To configure SNMPv3 notifications, follow these steps:

Procedure

Step 1 Select v3 from the Security drop-down list in the Create Destinations dialog box.

Step 2 (Optional) Set the inform time out and retry values.

Step 3 Click Create to add this destination to the selected switches.

Note In the case of SNMPv3 notifications, the SNMP manager is expected to know the user credentials (authKey/PrivKey) based on the switch's engineID to authenticate and decrypt the SNMP messages.

Enabling SNMP Notifications

This lists the DCNM-SAN procedures that enable the notifications for Cisco NX-OS MIBs. Expand Events > SNMP Traps to see the check boxes listed in this table.



Note Choosing Events > SNMP Traps enables both traps and informs, depending on how you configured SNMP notifications. See the notifications displayed with the *Configuring SNMPv3 Notifications*.

Table 44: Enabling SNMP Notifications

MIB	DCNM-SAN Check Boxes
CISCO-ENTITY-FRU-CONTROL-MIB	Click the Other tab and check FRU Changes.
CISCO-FCC-MIB	Click the Other tab and check FCC.
CISCO-DM-MIB	Click the FC tab and check Domain Mgr RCF.
CISCO-NS-MIB	Click the FC tab and check Name Server.
CISCO-FCS-MIB	Click the Other tab and check FCS Rejects.
CISCO-FDMI-MIB	Click the Other tab and check FDMI.
CISCO-FSPF-MIB	Click the FC tab and check FSPF Neighbor Change.
CISCO-LICENSE-MGR-MIB	Click the Other tab and check License Manager.
CISCO-IPSEC-SIGNALLING-MIB	Click the Other tab and check IPSEC.
CISCO-PSM-MIB	Click the Other tab and check Port Security.
CISCO-RSCN-MIB	Click the FC tab and check RSCN ILS, and RCSN ELS.
SNMPv2-MIB	Click the Other tab and check SNMP AuthFailure.
VRRP-MIB, CISCO-IETF-VRRP-MIB	Click the Other tab and check VRRP.

MIB	DCNM-SAN Check Boxes
CISCO-ZS-MIB	Click the FC tab and check Zone Rejects, Zone Merge Failures, Zone Merge Successes, Zone Default Policy Change, and Zone Unsuppd Mode.

The following notifications are enabled by default:

- entity fru
- license
- link ietf-extended

All other notifications are disabled by default.

You can enable or disable the supported traps at the following levels:

- Switch level—You can use `snmp-server enable traps` command to enable all the traps in the supported MIBs at the switch level.
- Feature level—You can use `snmp-server enable traps` command with the feature name to enable traps at the feature level.

```
switch =>snmp-server enable traps callhome ?
event-notify Callhome External Event Notification
smtp-send-fail SMTP Message Send Fail notification
```

- Individual traps - You can use `snmp-server enable traps` command with the feature name to enable traps at the individual level.

```
switch =>snmp-server enable traps callhome event-notify ?
```



Note

The `snmp-server enable traps` CLI command enables both traps and informs, depending on how you configured SNMP. See the notifications displayed with the `snmp-server host` CLI command.

1. Expand Events and then select SNMP Traps in the Physical Attributes pane.
You see the SNMP notification configuration in the Information pane.
2. Click the FC tab to enable Fibre Channel related notifications.
3. Check each notification check box that you want to enable.
4. Click the Other tab to enable other notifications.
5. Check each notification check box that you want to enable.
6. Click the Control tab to enable notification applicable variables.
7. From NX-OS Release 4.2(1), the Control tab is available for the notification control feature. This feature allows you to enable or disable all the notification-applicable variables via SNMP.
The Control tab is available for NX-OS Release 4.2(1) and later only.
8. Check each notification check box that you want to enable.

9. Click the Apply changes icon to create the entry.



Note In Device Manager, the `no snmp-server enable traps link` command disables generation of link traps in the switch, however the individual interfaces may have the link trap enabled.

To enable individual notifications, follow these steps:

1. Expand Admin > Events and then select Filters.

You see the event filters window showing a table populated by the switch

2. Click the Control tab to enable notification applicable variables.

From NX-OS Release 4.2(1), the Control tab is available for the notification control feature. This feature allows you to enable or disable all the notification-applicable variables via SNMP.



Note The Control tab is available for NX-OS Release 4.2(1) and later only.

3. Check each notification check box that you want to enable.
4. Click the Apply changes icon to create the entry.

Enable individual notifications

To enable individual notifications, follow these steps:

Procedure

- Step 1** Expand Events and then select SNMP Traps in the Physical Attributes pane.
You see the SNMP notification configuration in the Information pane.
- Step 2** Click the FC tab to enable Fibre Channel related notifications.
- Step 3** Check each notification check box that you want to enable.
- Step 4** Click the Other tab to enable other notifications.
- Step 5** Check each notification check box that you want to enable.
- Step 6** Click the Control tab to enable notification applicable variables.
- Step 7** From NX-OS Release 4.2(1), the Control tab is available for the notification control feature. This feature allows you to enable or disable all the notification-applicable variables via SNMP.

The Control tab is available for NX-OS Release 4.2(1) and later only.
- Step 8** Check each notification check box that you want to enable.
- Step 9** Click the Apply changes icon to create the entry.

Enable individual notifications using Device Manager

In Device Manager, the `no snmp-server enable traps link` command disables generation of link traps in the switch, however the individual interfaces may have the link trap enabled.

To enable individual notifications using Device Manager, follow these steps:

Procedure

-
- Step 1** Expand Admin > Events and then select Filters.
- You see the event filters window showing a table populated by the switch.
- Step 2** Click the Control tab to enable notification applicable variables.
- From NX-OS Release 4.2(1), the Control tab is available for the notification control feature. This feature allows you to enable or disable all the notification-applicable variables via SNMP.
- Note** The Control tab is available for NX-OS Release 4.2(1) and later only.
- Step 3** Check each notification check box that you want to enable.
- Step 4** Click the Apply changes icon to create the entry.
-

Configuring the Notification Target User

You must configure a notification target user on the switch for sending SNMPv3 inform notifications to the SNMP manager.

For authenticating and decrypting the received INFORM PDU, the SNMP manager should have the same user credentials in its local configuration data store of users.

To configure the notification target user, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

Configuring LinkUp/LinkDown Notifications for Switches

To configure the LinkUp/LinkDown notification for a switch using NX-OS Release 4.1(x) and earlier, follow these steps:



-
- Note** If both IETF and IETF extended are enabled, the `show snmp traps` command displays both as enabled. However, as a trap, you will receive only one trap with IETF extended payload.
-

Configuring Up/Down SNMP Link-State Traps for Interfaces

By default, SNMP link-state traps are enabled for all interfaces. Whenever a link toggles its state from Up to Down or vice versa, an SNMP trap is generated.

In some instances, you may find that you have numerous switches with hundreds of interfaces, many of which do not require monitoring of the link state. In such cases, you may elect to disable link-state traps.

Configuring Entity (FRU) Traps



Note All these traps have to do with legacy FRU traps.

Configuring Event Security

SNMP events can be secured against interception or eavesdropping in the same way that SNMP messages are secured. DCNM-SAN or Device Manager allow you to configure the message processing model, the security model, and the security level for the SNMP events that the switch generates.



Note This is an advanced function that should only be used by administrators having experience with SNMPv3.

To configure SNMP event security, follow these steps:

Procedure

- Step 1** Expand Events and then select SNMP Traps.
- Step 2** Click the Security tab in the Information pane.
You see the security information for SNMP notifications.
- Step 3** Set the message protocol model (MPModel), security model, security name, and security level.
- Step 4** Click the Apply Changes icon to save and apply your changes.

Viewing the SNMP Events Log

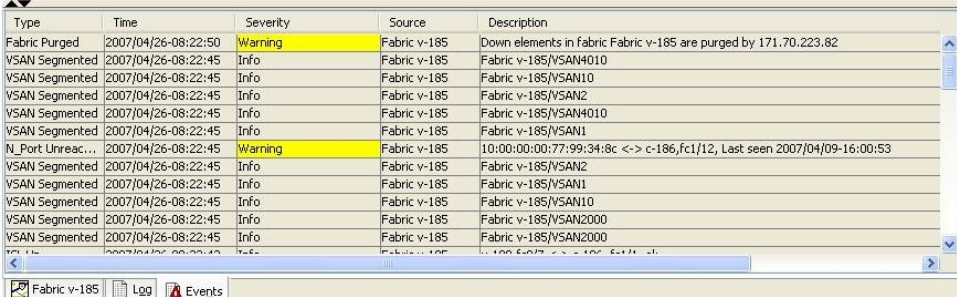
- You must set up the MDS syslog manager before you can view the event logs.
- Changing these values from different DCNM-SAN workstations at the same time may cause unpredictable results.

Procedure

To view the SNMP events log from DCNM-SAN, click the Events tab.

You see the Events listed with a log of events for a single switch in the following image.

Figure 42: Events Information



Type	Time	Severity	Source	Description
Fabric Purged	2007/04/26-08:22:50	Warning	Fabric v-185	Down elements in fabric Fabric v-185 are purged by 171.70.223.82
YSAN Segmented	2007/04/26-08:22:45	Info	Fabric v-185	Fabric v-185/YSAN4010
YSAN Segmented	2007/04/26-08:22:45	Info	Fabric v-185	Fabric v-185/YSAN10
YSAN Segmented	2007/04/26-08:22:45	Info	Fabric v-185	Fabric v-185/YSAN2
YSAN Segmented	2007/04/26-08:22:45	Info	Fabric v-185	Fabric v-185/YSAN4010
YSAN Segmented	2007/04/26-08:22:45	Info	Fabric v-185	Fabric v-185/YSAN1
N_Port Unreac...	2007/04/26-08:22:45	Warning	Fabric v-185	10:00:00:00:77:99:34:8c <-> c-186,fc1/12, Last seen 2007/04/09-16:00:53
YSAN Segmented	2007/04/26-08:22:45	Info	Fabric v-185	Fabric v-185/YSAN2
YSAN Segmented	2007/04/26-08:22:45	Info	Fabric v-185	Fabric v-185/YSAN1
YSAN Segmented	2007/04/26-08:22:45	Info	Fabric v-185	Fabric v-185/YSAN10
YSAN Segmented	2007/04/26-08:22:45	Info	Fabric v-185	Fabric v-185/YSAN2000
YSAN Segmented	2007/04/26-08:22:45	Info	Fabric v-185	Fabric v-185/YSAN2000

Field Descriptions for SNMP

This section describes the field descriptions for SNMP.

IP Statistics SNMP

Field	Description
BadVersions	The total number of SNMP messages which were delivered to the SNMP entity and were for an unsupported SNMP version.
BadCommunityNames	The total number of SNMP messages delivered to the SNMP entity which used a SNMP community name not known to said entity.
BadCommunityUses	The total number of SNMP messages delivered to the SNMP entity which represented an SNMP operation which was not allowed by the SNMP community named in the message.
ASNParseErrs	The total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages.
TooBigs	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is tooBig.
SilentDrops	The total number of GetRequest-PDUs, GetNextRequest-PDUs, GetBulkRequest-PDUs, SetRequest-PDUs, and InformRequest-PDUs delivered to the SNMP entity which were silently dropped because the size of a reply containing an alternate Response-PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the request.
ProxyDrops	The total number of GetRequest-PDUs, GetNextRequest-PDUs, GetBulkRequest-PDUs, SetRequest-PDUs, and InformRequest-PDUs delivered to the SNMP entity which were silently dropped because the transmission of the (possibly translated) message to a proxy target failed in a manner (other than a time-out) such that no Response-PDU could be returned.

Field	Description
NoSuchNames	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is noSuchName.
BadValues	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is badValue.
ReadOnlys	The total number valid SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is readOnly. It should be noted that it is a protocol error to generate an SNMP PDU which contains the value readOnly in the error-status field, as such this is provided as a means of detecting incorrect implementations of the SNMP.
GenErrs	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is genErr.
Pkts	The total number of messages delivered to the SNMP entity from the transport service.
GetRequests	The total number of SNMP Get-Request PDUs which have been accepted and processed by the SNMP protocol entity.
GetNexts	The total number of SNMP Get-Next PDUs which have been accepted and processed by the SNMP protocol entity.
SetRequests	The total number of SNMP Set-Request PDUs which have been accepted and processed by the SNMP protocol entity.
OutTraps	The total number of SNMP Trap PDUs which have been generated by the SNMP protocol entity.
OutGetResponses	The total number of SNMP Get-Response PDUs which have been generated by the SNMP protocol entity.
OutPkts	The total number of SNMP Messages which were passed from the SNMP protocol entity to the transport service.
TotalReqVars	The total number of MIB objects which have been retrieved successfully by the SNMP protocol entity as the result of receiving valid SNMP Get-Request and Get-Next PDUs.
TotalSetVars	The total number of MIB objects which have been altered successfully by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs.

SNMP Security Users

Field	Description
Role	The user in Security Model independent format.
Password	Password of the common user. For SNMP, this password is used for both authentication and privacy. For CLI and XML, it is used for authentication only.

Field	Description
Digest	The type of digest authentication protocol which is used.
Encryption	The type of encryption authentication protocol which is used.
ExpiryDate	The date on which this user will expire.
SSH Key File Configured	Specifies whether the user is configured with SSH public key.
SSH Key File	The name of the file storing the SSH public key. The SSH public key is used to authenticate the SSH session for this user. Note that this applies to only CLI user. The format can be one of the following: <ul style="list-style-type: none"> • SSH Public Key in OpenSSH format • SSH Public Key in IETF SECSH (Commercial SSH public key format) • SSH Client Certificate in PEM (privacy-enhanced mail format) from which the public key is extracted • SSH Client Certificate DN (Distinguished Name) for certificate based authentication
Creation Type	The type of the credential store of the user. When a row is created in this table by a user, the user entry is created in a credential store local to the device. In case of remote authentication mechanism like AAA Server based authentication, credentials are stored in other (remote) system/device.
Expiry Date	The date on which this user will expire.

SNMP Security Communities

Field	Description
Community	The community string.
Role	The Security Model name.

For more information, refer *Adding or Deleting Communities* and *Deleting a Community String*.

Security Users Global

Field	Description
Enforce SNMP Privacy Encryption	Specifies whether the SNMP agent enforces the use of encryption for SNMPv3 messages globally on all the users in the system.
Cache Timeout	This specifies maximum timeout value for caching the user credentials in the local system.



Note The privacy password and authentication password are required for an administrator to create a new user or delete an existing user in Device Manager. However, if the administrator does not provide these credentials at the time of creating a new user, Device Manager uses the authentication password of the administrator as the privacy password. If the privacy protocol defined for the user is not DES (default), the SNMP Agent in the MDS will not be able to decrypt the packet and the SNMP Agent times out. If the privacy protocol defined for the user is not DES, the user needs to provide both the privacy password and the protocol when logging in.



CHAPTER 15

Configuring Domain Parameters

- [Configuring Domain Parameters, on page 383](#)

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.

This chapter includes the following sections:

Information About Fibre Channel Domains

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.

This section describes each fcdomain phase:

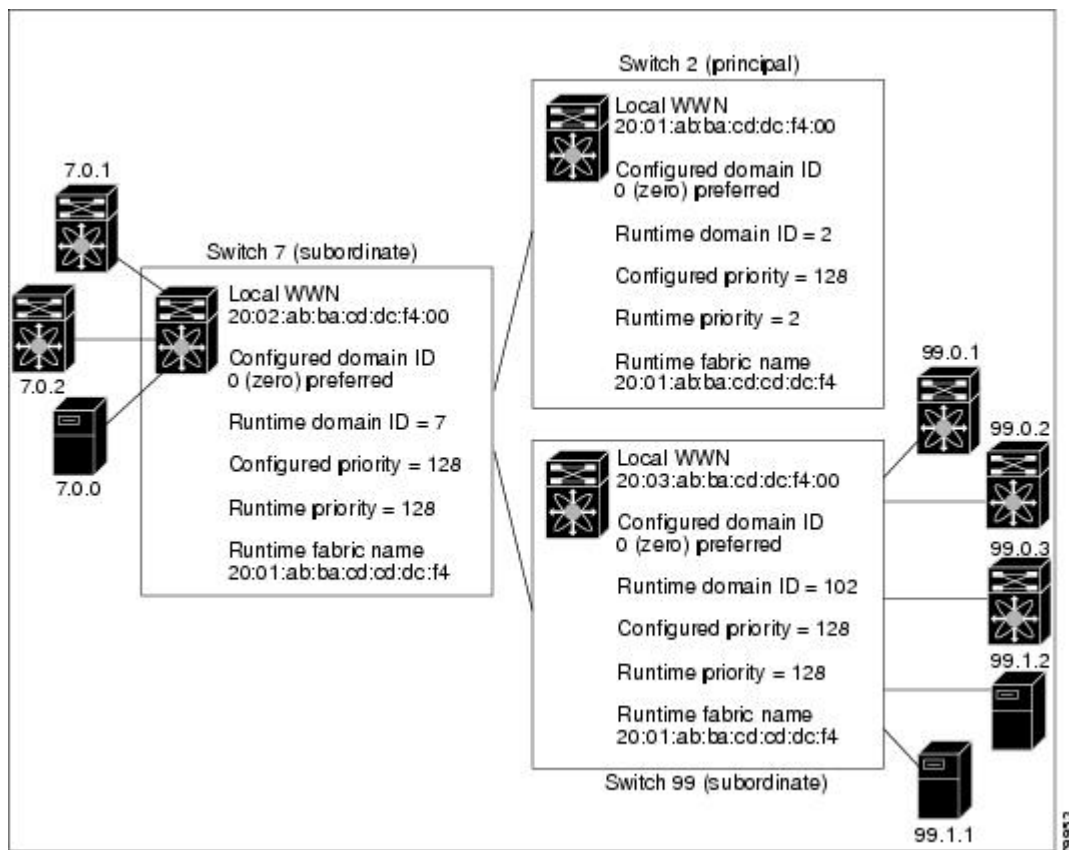
- Principal switch selection—This phase guarantees the selection of a unique principal switch across the fabric.
- Domain ID distribution—This phase guarantees each switch in the fabric obtains a unique domain ID.
- FC ID allocation—This phase guarantees a unique FC ID assignment to each device attached 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



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.

The figure below shows a sample fcdomain configuration.



This section includes the following topics:

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 restart the fcdomain for that VSAN disruptively.

You can apply most of the configurations to their corresponding runtime values. Each of the following sections provide further details on how the `fcdomain` parameters are applied to the runtime values.

The **`fcdomain restart`** command applies your changes to the runtime settings. Use the **`disruptive`** option to apply most of the configurations to their corresponding runtime values, including preferred domain IDs (see the [Domain IDs, on page 386](#)).

Domain Manager Fast Restart

As of Cisco MDS SAN-OS Release 3.0(2), when a principal link fails, the domain manager must select a new principal link. By default, the domain manager starts a build fabric (BF) phase, followed by a principal switch selection phase. Both of these phases involve all the switches in the VSAN and together take at least 15 seconds to complete. To reduce the time required for the domain manager to select a new principal link, you can enable the domain manager fast restart feature.

When fast restart is enabled and a backup link is available, the domain manager needs only a few milliseconds to select a new principal link to replace the one that failed. Also, the reconfiguration required to select the new principal link only affects the two switches that are directly attached to the failed link, not the entire VSAN. When a backup link is not available, the domain manager reverts to the default behavior and starts a BF phase, followed by a principal switch selection phase. The fast restart feature can be used in any interoperability mode.

**Tip**

We recommend using fast restart on most fabrics, especially those with a large number of logical ports (3200 or more), where a logical port is an instance of a physical port in a VSAN.

Switch Priority

Any new switch can become the principal switch when it joins a stable fabric. During the principal switch selection phase, the switch with the highest priority becomes the principal switch. If two switches have the same configured priority, the switch with the lower WWN becomes the principal switch.

The priority configuration is applied to runtime when the `fcdomain` is restarted (see the [Domain Restart, on page 384](#)). This configuration is applicable to both disruptive and nondisruptive restarts.

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.

Incoming RCFs

You can choose to reject RCF request frames on a per-interface, per-VSAN basis. By default, the RCF reject option is disabled (that is, RCF request frames are not automatically rejected).

The RCF reject option takes immediate effect at runtime through a disruptive restart (see the [Domain Restart, on page 384](#)).

You can configure the `rcf-reject` option 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 is required.

Autoreconfiguring Merged Fabrics

By default, the autoreconfigure option is disabled. When you join two switches belonging to two different stable fabrics that have overlapping domains, the following cases apply:

- If the autoreconfigure option is enabled on both switches, a disruptive reconfiguration phase is started.
- If the autoreconfigure option is disabled on either or both switches, the links between the two switches become isolated.

The autoreconfigure option takes immediate effect at runtime. You do not need to restart the fcdomain. If a domain is currently isolated due to domain overlap, and you later enable the 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.

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.

The configured domain ID can be preferred or static. By default, the configured domain ID is 0 (zero) and the configured type is preferred.



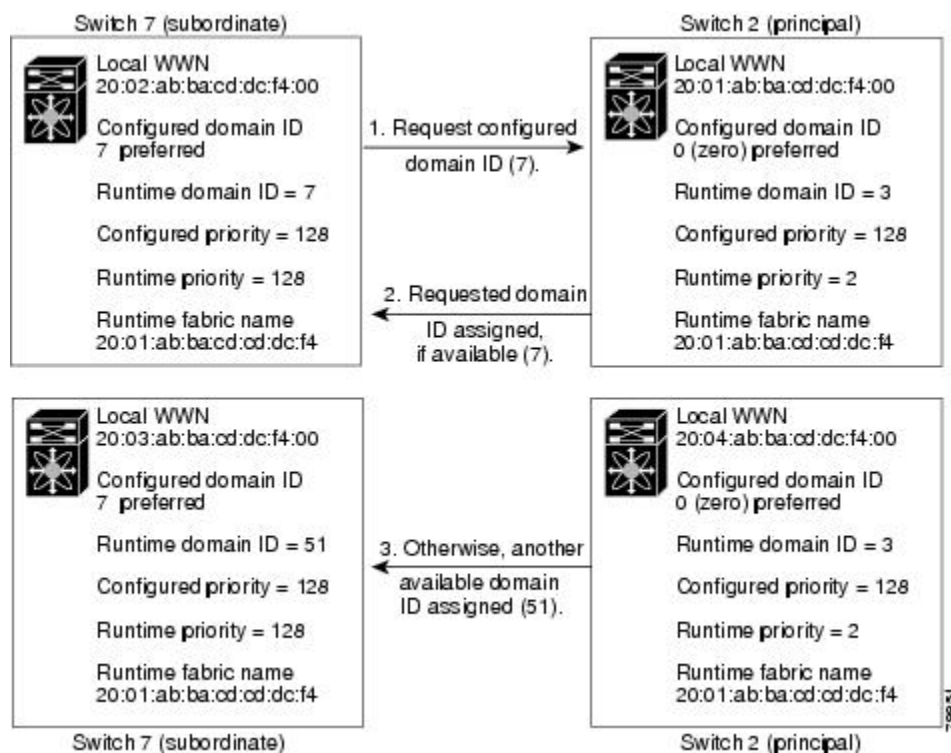
Note

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:

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.



The behavior for a subordinate switch changes based on three factors:

- The allowed domain ID lists.
- The configured domain ID.
- The domain ID that the principal switch has assigned to the requesting switch.

In specific situations, the changes are as follows:

- When the received domain ID is not within the allowed list, the requested domain ID becomes the runtime domain ID and all interfaces on that VSAN are isolated.
- When the assigned and requested domain IDs are the same, the preferred and static options are not relevant, and the assigned domain ID becomes the runtime domain ID.
- When the assigned and requested domain IDs are different, the following cases apply:
 - If the configured type is static, the assigned domain ID is discarded, all local interfaces are isolated, and the local switch assigns itself the configured domain ID, which becomes the runtime domain ID.
 - If the configured type is preferred, the local switch accepts the domain ID assigned by the principal switch and the assigned domain ID becomes the runtime domain ID.

If you change the configured domain ID, the change is only accepted if the new domain ID is included in all the allowed domain ID lists currently configured in the VSAN. Alternatively, you can also configure zero-preferred domain ID.

**Tip**

When the FICON feature is enabled in a given VSAN, the domain ID for that VSAN remains in the static state. You can change the static ID value but you cannot change it to the preferred option.

**Note**

In an IVR without NAT configuration, if one VSAN in the IVR topology is configured with static domain IDs, then the other VSANs (edge or transit) in the topology should also be configured with static domain IDs. In an IVR NAT configuration, if one VSAN in the IVR topology is configured with static domain IDs, then the IVR domains that can be exported to that VSAN must also be assigned static domains.

**Caution**

You must enter the `fcdomain restart` command if you want to apply the configured domain changes to the runtime domain.

**Caution**

You must restart the **fcdomain** if you want to apply the configured domain changes to the runtime domain.

**Note**

If you have configured an allowed domain ID list, the domain IDs that you add must be in that range for the VSAN. See the *Configuring Allowed Domain ID Lists* section.

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 get 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 [Domain Restart, on page 384](#)).

Allowed Domain ID Lists

By default, the valid range for an assigned domain ID list is from 1 to 239. You can specify a list of ranges to be in the allowed domain ID list and separate each range with a comma. The principal switch assigns domain IDs that are available in the locally configured allowed domain list.

Use allowed domain ID lists to design your VSANs with non-overlapping domain IDs. This helps you in the future if you need to implement IVR without the NAT feature.

CFS Distribution of Allowed Domain ID Lists

You can enable the distribution of the allowed domain ID lists configuration information to all Cisco MDS switches in the fabric using the Cisco Fabric Services (CFS) infrastructure. This feature allows you to synchronize the configuration across the fabric from the console of a single MDS switch. Since the same configuration is distributed to the entire VSAN, you avoid possible misconfiguration and the likelihood that two switches in the same VSAN have configured incompatible allowed domains.

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 allow domain ID list and committing it on the principle switch.

For more information about CFS, see *Chapter 13, “Using the CFS Infrastructure.”*

Contiguous Domain ID Assignments

By default, the contiguous domain assignment is disabled. When a subordinate switch requests the principal switch for two or more domains and the domains are not contiguous, the following cases apply:

- If the contiguous domain assignment is enabled in the principal switch, the principal switch locates contiguous domains and assigns them to the subordinate switches. If contiguous domains are not available, the NX-OS software rejects this request.
- If the contiguous domain assignment is disabled in the principal switch, the principal switch assigns the available domains to the subordinate switch.

Locking the Fabric

The first action that modifies the existing configuration creates the pending configuration and locks the feature in the fabric. Once you lock the fabric, the following conditions apply:

- No other user can make any configuration changes to this feature.
- A pending configuration is created by copying the active configuration. Modifications from this point on are made to the pending configuration and remain there until you commit the changes to the active configuration (and other switches in the fabric) or discard them.

Committing Changes

To apply the pending domain configuration changes to other MDS switches in the VSAN, you must commit the changes. The pending configuration changes are distributed and, on a successful commit, the configuration changes are applied to the active configuration in the MDS switches throughout the VSAN and the fabric lock is released.

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.

The pending changes are only available in the volatile directory and are discarded if the switch is restarted.

FC IDs

When an N or NL port logs into a Cisco MDS 9000 Family switch, it is assigned an FC ID. By default, the persistent FC ID feature is enabled. If this feature is disabled, the following consequences apply:

- An N or NL port logs into a Cisco MDS 9000 Family switch. The WWN of the requesting N or NL port and the assigned FC ID are retained and stored in a volatile cache. The contents of this volatile cache are not saved across reboots.

- The switch is designed to preserve the binding FC ID to the WWN on a best-effort basis. For example, if one N port disconnects from the switch and its FC ID is requested by another device, this request is granted and the WWN with the initial FC ID association is released.
- The volatile cache stores up to 4000 entries of WWN to FC ID binding. If this cache is full, a new (more recent) entry overwrites the oldest entry in the cache. In this case, the corresponding WWN to FC ID association for the oldest entry is lost.
- The switch connection behavior differs between N ports and NL ports:
 - N ports receive the same FC IDs if disconnected and reconnected to any port within the same switch (as long as it belongs to the same VSAN).
 - NL ports receive the same FC IDs only if connected back to the same port on the switch to which they were originally connected.

Persistent FC IDs

When persistent FC IDs are enabled, the following consequences apply:

- The currently *in use* FC IDs in the fcdomain are saved across reboots.
- The fcdomain automatically populates the database with dynamic entries that the switch has learned about after a device (host or disk) is plugged into a port interface.

Persistent FC ID Configuration

When the persistent FC ID feature is enabled, you can enter the persistent FC ID submode and add static or dynamic entries in the FC ID database. By default, all added entries are static. Persistent FC IDs are configured on a per-VSAN basis. Follow these requirements to manually configure a persistent FC ID:

- Ensure that the persistent FC ID feature is enabled in the required VSAN.
- Ensure that the required VSAN is an active VSAN—persistent FC IDs can only be configured on active VSANs.
- Verify that the domain part of the FC ID is the same as the runtime domain ID in the required VSAN. If the software detects a domain mismatch, the command is rejected.
- Verify that the port field of the FC ID is 0 (zero) when configuring an area.



Note FICON uses a different scheme for allocating FC IDs based in the front panel port number. This scheme takes precedence over FC ID persistence in FICON VSANs.

About Unique Area FC IDs for HBAs



Note Read this section only if the HBA port and the storage port are connected to the same switch.

Some HBA ports require a different area ID than storage ports when they are both connected to the same switch. For example, if the storage port FC ID is 0x6f7704, the area for this port is 77. In this case, the HBA port's area can be anything other than 77. The HBA port's FC ID must be manually configured to be different from the storage port's FC ID.

Switches in the Cisco MDS 9000 Family facilitate this requirement with the FC ID persistence feature. You can use this feature to preassign an FC ID with a different area to either the storage port or the HBA port.

Persistent FC ID Selective Purging

Persistent FC IDs can be purged selectively. Static entries and FC IDs currently in use cannot be deleted. [Table 45: Purged FC IDs](#), on page 391 identifies the FC ID entries that are deleted or retained when persistent FC IDs are purged.

Table 45: 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

Guidelines and Limitations

- 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.
- 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.

Default Settings

[Table 46: Default fcdomain Parameters](#), on page 391 lists the default settings for all fcdomain parameters.

Table 46: Default fcdomain Parameters

Parameters	Default
fcdomain feature	Enabled.
Configured domain ID	0 (zero).
Configured domain	Preferred.
auto-reconfigure autoreconfigureoption	Disabled.
contiguous-allocation option	Disabled.
Priority	128.
Allowed list	1 to 239.
Fabric name	20:01:00:05:30:00:28:df.

Parameters	Default
rcf-reject	Disabled.
Persistent FC ID	Enabled.
Allowed domain ID list configuration distribution	Disabled.

Configuring Fibre Channel Domains

This section describes the fcdomain feature and includes the following topics:

Configuring Domain Manager Turbo Mode

The Domain Manager turbo mode feature allows you to restart the Domain Manager with optimization. You have the option to select fast-restart or selective-restart mode for restarting the Domain Manager. You can leave the restart mode empty indicating that optimization is disabled.

Configuring Domain Manager Turbo Mode

To configure the Domain Manager turbo mode, follow these steps:

Procedure

-
- Step 1** Expand Fabric > All VSANs and then select Domain Manager in the Logical Domains pane for the fabric and VSAN for which you want to configure turbo mode. You see the Running tab configuration of the domain in the Information pane.
 - Step 2** Click the Configuration tab.
 - Step 3** Set the Optimization drop-down menu to fast-restart or selective-restart for any switch in the fabric that you want to optimize. You can leave the Optimization field without any selection, indicating that the optimization is disabled.
 - Step 4** Click the Apply Changes icon to initiate this restart.
-

Configuring the Domain Manager Turbo Mode Using Device Manager

To configure the Domain Manager turbo mode using Device Manager, follow these steps:

Procedure

-
- Step 1** Expand FC > Domain Manager and then select the Configuration tab.
Note The Optimization field is not available in releases prior to NX-OS Release 4.2(1).
 - Step 2** Set the Optimization drop-down menu to fast-restart or selective-restart for any switch in the fabric that you want to optimize. You can leave the Optimization field without any selection, indicating that the optimization is disabled.

- Step 3** Click Apply to initiate this restart.
-

Restarting a Domain

Domain Configuration Scenarios:

- Switch Configuration: Irrespective of how the switches in VSAN 6 are configured, fcdomain restart disruptive vsan 6 causes all devices of all switches in VSAN 6 to log out, causing data traffic disruption.
- Configured domain and the runtime domain are the same: Assuming that the configured domain and the runtime domain are the same on all switches, fcdomain restart vsan 6 does not cause any devices in VSAN 6 to log out.
- Configured domain and runtime domain are not the same: Assuming that on some switches in VSAN 6 the configured domain and the runtime domain are not the same, fcdomain restart vsan 6 causes the devices in VSAN 6 attached to the switches whose statically configured and runtime domain differ to log out, causing data traffic disruption.

Restarting a Domain

To restart the fabric disruptively or nondisruptively, follow these steps:

Procedure

- Step 1** Expand Fabric > All VSANs and then select Domain Manager in the Logical Domains pane for the fabric and VSAN that you want to restart.
- Step 2** Click the **Configuration** tab.
- Step 3** Set the Restart drop-down menu to disruptive or nonDisruptive for any switch in the fabric that you want to restart the fcdomain.
- Step 4** Click the Apply Changes icon to initiate this fcdomain restart.
-

Configuring 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.

To configure the priority for the principal switch, follow these steps:

Procedure

- Step 1** Expand Fabric > All VSANs and then select Domain Manager in the Logical Domains pane for the fabric and VSAN that you want to set the principal switch priority for.
- Step 2** Set Priority to a high value for the switch in the fabric that you want to be the principal switch.
- Step 3** Click the Apply Changes icon to save these changes.
-

Enabling or Disabling fcdomains

To disable or reenable fcdomains in a single VSAN or a range of VSANs, follow these steps:

Procedure

- Step 1** Expand Fabric > All VSANs and then select Domain Manager in the Logical Domains pane for the fabric and VSAN that you want to disable fcdomain for.
- You see the domain's running configuration in the Information pane.
- Step 2** Click the Configuration tab and uncheck the Enable check box for each switch in the fabric that you want to disable fcdomain on.
- Step 3** Click the Apply Changes icon to save these changes.
-

Configuring Fabric Names

To set the fabric name value for a disabled fcdomain, follow these steps:

Procedure

- Step 1** Expand Fabric > All VSANs and then select Domain Manager in the Logical Domains pane for the fabric and VSAN that you want to set the fabric name for.
- You see the running configuration of the domain in the Information pane.
- Step 2** Click the Configuration tab and set the fabric name for each switch in the fabric.
- Step 3** Click the Apply Changes icon to save these changes.
-

Rejecting Incoming RCFs

To reject incoming RCF request frames, follow these steps:

Procedure

- Step 1** Expand Switches > FC Interfaces and then select Physical in the Physical Attributes pane.
- You see the Fibre Channel configuration in the Information pane.
- Step 2** Click the Domain Mgr tab.
- 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.
-

Enabling Autoreconfiguration

To enable automatic reconfiguration in a specific VSAN (or range of VSANs), follow these steps:

Procedure

-
- | | |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Expand Fabric > All VSANs and then select Domain Manager in the Logical Domains pane for the fabric and VSAN that you want to enable automatic reconfiguration for.

You see the running configuration of the domain in the Information pane. |
| Step 2 | Select the Configuration tab and check the Auto Reconfigure check box for each switch in the fabric that you want to automatically reconfigure. |
| Step 3 | Click the Apply Changes icon to save these changes. |
-

Configuring 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.

The configured domain ID can be preferred or static. By default, the configured domain ID is 0 (zero) and the configured type is preferred.

This section includes the following topics:

Specifying Static or Preferred Domain IDs

When a new domain ID is configured, the new configuration has to be applied by manually restarting the domain using the `fcdomain restart` command; if a discrepancy is detected between the configured domain ID and the runtime domain ID during the subsequent fabric merge, the link will be isolated.

To specify a static or preferred domain ID, follow these steps:

Procedure

-
- | | |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Expand Fabric > All VSANs and then select Domain Manager in the Logical Domains pane for the fabric and VSAN that you want to configure the domain ID for.

You see the running configuration of the domain in the Information pane. |
| Step 2 | Enter a value for the Config DomainID and click static or preferred from the Config Type drop-down menu to set the domain ID for switches in the fabric. |
| Step 3 | Click the Apply Changes icon to save these changes. |
-

Configuring Allowed Domain ID Lists

An allowed domain ID list must satisfy the following conditions:

- If this switch is a principal switch, all the currently assigned domain IDs must be in the allowed list.

- If this switch is a subordinate switch, the local runtime domain ID must be in the allowed list.
- The locally configured domain ID of the switch must be in the allowed list.
- The intersection of the assigned domain IDs with other already configured domain ID lists must not be empty.

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.

To configure the allowed domain ID list, follow these steps:

Procedure

-
- Step 1** Expand Fabric > All VSANs > Domain Manager and then select Allowed in the Logical Domains pane for the fabric and VSAN for which you want to set the allowed domain ID list.
- You see the CFS configuration in the Information pane.
- Step 2** Set the Admin drop-down menu to enable and set the Global drop-down menu to enable.
- Step 3** Click Apply Changes to enable CFS distribution for the allowed domain ID list.
- Step 4** Select the Allowed DomainIds tab.
- Step 5** Set the list to the allowed domain IDs list for this domain.
- Step 6** Select the CFS tab and set Config Action to commit.
- Step 7** Click the Apply Changes icon to commit this allowed domain ID list and distribute it throughout the VSAN.
-

Enabling Allowed Domain ID 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, follow these steps:

Before you begin

All switches in the fabric must be running Cisco SAN-OS Release 3.0(1) or later to distribute the allowed domain ID list using CFS.

Procedure

-
- Step 1** Expand Fabric > All VSANs > Domain Manager and then select Allowed in the Logical Domains pane for the fabric and VSAN that you want to set the allowed domain ID list for.
- You see the CFS configuration in the Information pane.
- Step 2** Set the Admin drop-down menu to enable and the Global drop-down menu to enable to enable CFS distribution for the allowed domain ID list.
- Step 3** Click the Apply Changes icon to enable CFS distribution for the allowed domain ID list.
-

Committing Changes

To commit pending domain configuration changes and release the lock, follow these steps:

Procedure

-
- Step 1** Expand Fabric > All VSANs > Domain Manager and then select Allowed in the Logical Domains pane for the fabric and VSAN that you want to set the allowed domain ID list for.
- You see the CFS configuration in the Information pane.
- Step 2** Set the Config Action drop-down menu to commit.
- Step 3** Click 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, follow these steps:

Procedure

-
- Step 1** Expand Fabric > All VSANs > Domain Manager and then select Allowed in the Logical Domains pane for the fabric and VSAN that you want to set the allowed domain ID list for.
- You see the CFS configuration in the Information pane.
- Step 2** Set the Config Action drop-down menu to abort.
- Step 3** Click the Apply Changes icon to discard any pending changes to the allowed domain ID list.
-

Enabling Contiguous Domain ID Assignments

To enable contiguous domains in a specific VSAN (or a range of VSANs), follow these steps:

Procedure

-
- Step 1** Expand Fabric > All VSANs and then select Domain Manager in the Logical Domains pane for the fabric and VSAN that you want to enable contiguous domains for.
- You see the running configuration of the domain in the Information pane.
- Step 2** 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.
-

Configuring FC IDs

When an N or NL port logs into a Cisco MDS 9000 Family switch, it is assigned an FC ID.

This section includes the following topics:

Enabling the Persistent FC ID Feature

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. A persistent FC ID assigned to an F port can be moved across interfaces and can continue to maintain the same persistent FC ID.

- FC IDs are enabled by default. This change of default behavior from releases prior to Cisco MDS SAN-OS Release 2.0(1b) prevents FC IDs from being changed after a reboot. You can disable this option for each VSAN.
- Persistent FC IDs with loop-attached devices (FL ports) need to remain connected to the same port in which they were configured.
- Due to differences in Arbitrated Loop Physical Address (ALPA) support on devices, FC ID persistency for loop-attached devices is not guaranteed.

To enable the persistent FC ID feature, follow these steps:

Procedure

-
- Step 1** Expand Fabric > All VSANs and then select Domain Manager in the Logical Domains pane for the fabric and VSAN that you want to enable the Persistent FC ID feature for.
You see the running configuration of the domain in the Information pane.
 - Step 2** Select the Persistent Setup tab and check the enable check box for each switch in the fabric that will have persistent FC ID enabled.
 - Step 3** Click the Apply Changes icon to save these changes.
-

Configuring Persistent FC IDs

To configure persistent FC IDs, follow these steps:

Procedure

-
- Step 1** Expand Fabric > All VSANs and then select Domain Manager in the Logical Domains pane for the fabric and VSAN that you want to configure the Persistent FC ID list for.
You see the running configuration of the domain in the Information pane.
 - Step 2** Click the Persistent FcIds tab and click Create Row.
 - Step 3** Select the switch, WWN, and FC ID that you want to make persistent.
 - Step 4** Set the Mask radio button to single or area.
 - Step 5** Set the Assignment radio button to static or dynamic.

- Step 6** Click the Apply Changes icon to save these changes.
-

Configuring Unique Area FC IDs for an HBA

To configure a different area ID for the HBA port, follow these steps:

Procedure

- Step 1** Expand End Device in the Physical Attributes pane and select the FLOGI tab in the Information pane to obtain the port WWN (Port Name field) of the HBA.
- Note** Both FC IDs in this setup have the same area 00 assignment.
- Step 2** Expand Switches > FC Interfaces and then select Physical from the Physical Attributes pane.
- Step 3** Set the Status Admin drop-down menu to down for the interface that the HBA is connected to. This shuts down the HBA interface in the MDS switch.
- Step 4** Expand Fabric > All VSANs and then select Domain Manager.
- Step 5** Click the Persistent Setup tab in the Information pane to verify that the FC ID feature is enabled. If this feature is disabled, continue with this procedure to enable the persistent FC ID. If this feature is already enabled, skip to Step 7.
- Step 6** Check the **Enable** check box to enable the persistent FC ID feature in the Cisco MDS switch.
- Step 7** Select the Persistent FeIds tab and assign a new FC ID with a different area allocation in the FeId field. In this example, we replace 00 with ee.
- Step 8** Click Apply Changes to save this new FC ID.
- Step 9** Compare the FC ID values to verify the FC ID of the HBA.
- Note** Both FC IDs now have different area assignments.
- Step 10** Expand Switches > FC Interfaces and then select Physical from the Physical Attributes pane. Set the Status Admin drop-down menu to up for the interface that the HBA is connected to. This enables the HBA interface in the MDS switch.
-

Purging Persistent FC IDs

To purge persistent FC IDs, follow these steps:

Procedure

- Step 1** Expand Fabric > 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.
- Step 3** Check the Purge check box for the switch that you want to purge persistent FC IDs on.

- Step 4** Click the Apply Changes icon to save these changes.
-

Clearing a Fabric Lock

To release a fabric lock, follow these steps:

Procedure

- Step 1** Expand Fabric > All VSANs > Domain Manager and then select AllowedId in the Logical Domains pane for the fabric and VSAN for which you want the allowed domain ID list.
- You see the CFS configuration in the Information pane.
- Step 2** Set the Config Action drop-down menu to clear.
- Step 3** Click the Apply Changes icon to clear the fabric lock.
-

Displaying Pending Changes

To display the pending configuration changes, follow these steps:

Procedure

- Step 1** Expand Fabric > All VSANs > Domain Manager > Allowed in the Logical Domains pane for the fabric and VSAN that you want to set the allowed domain ID list for.
- You see the CFS configuration in the Information pane.
- Step 2** Set the Config View As drop-down menu to pending.
- Step 3** Click 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, follow these steps:

Procedure

- Step 1** Expand Fabric > All VSANs > Domain Manager and then select Allowed in the Logical Domains pane for the fabric and VSAN for which you want to set the allowed domain ID list.
- Step 2** View the CFS configuration and session status in the Information pane.
-

Monitoring FC Domain

This section covers the following topic:

Displaying fcdomain Statistics

To display fcdomain statistics, follow these steps:

Procedure

-
- Step 1** Expand Fabric > All VSANs and then select Domain Manager in the Logical Domains pane for the fabric that you want to display statistics for.
- You see the running configuration of the domain in the Information pane.
- Step 2** Click the Statistics tab. You see the FC ID statistics in the Information pane.
-

Field Descriptions for FC Domain

This section describes the field descriptions for FC Domain.

IVR Domains

Field	Description
Domain Id	The FC domain ID that will be used to represent the VSAN.



CHAPTER 16

Configuring and Managing Zones

- [Configuring and Managing Zones, on page 403](#)

Configuring and Managing Zones

Zoning enables you to set up access control between storage devices or user groups. If you have administrator privileges in your fabric, you can create zones to increase network security and to prevent data loss or corruption. Zoning is enforced by examining the source-destination ID field.

Advanced zoning capabilities specified in the FC-GS-4 and FC-SW-3 standards are provided. You can use either the existing basic zoning capabilities or the advanced, standards-compliant zoning capabilities.

For information about design parameters and best practices to Migrate a SAN from a Heterogeneous Environment to a Cisco MDS 9000 Family SAN, refer to [Migrate a SAN from a Heterogeneous Environment to a Cisco MDS 9000 Family SAN](#).

This chapter includes the following topics:

Information About Zoning

Zoning has the following features:

- A zone consists of multiple zone members:
 - Members in a zone can access each other; members in different zones cannot access each other.
 - If zoning is not activated, all devices are members of the default zone.
 - If zoning is activated, any device that is not in an active zone (a zone that is part of an active zone set) is a member of the default zone.
 - Zones can vary in size.
 - Devices can belong to more than one zone.
 - A 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 criteria is based mainly on WWNs or FC IDs.
 - Port world wide name (pWWN)—Specifies the pWWN of an N port attached to the switch as a member of the zone.
 - Fabric pWWN—Specifies the WWN of the fabric port (switch port's WWN). This membership is also referred to as port-based zoning.
 - FC ID—Specifies the FC ID of an N port attached to the switch as a member of the zone.
 - Interface and switch WWN (sWWN)—Specifies the interface of a switch identified by the sWWN. This membership is also referred to as interface-based zoning.
 - Interface and domain ID—Specifies the interface of a switch identified by the domain ID.
 - Domain ID and port number—Specifies the domain ID of an MDS domain and additionally specifies a port belonging to a non-Cisco switch.
 - IPv4 address—Specifies the IPv4 address (and optionally the subnet mask) of an attached device.
 - IPv6 address—The IPv6 address of an attached device in 128 bits in colon(:)-separated hexadecimal format.
- Default zone membership includes all ports or WWNs that do not have a specific membership association. Access between default zone members is controlled by the default zone policy.
- You can configure up to 8000 zones per VSAN and a maximum of 8000 zones for all VSANs on the switch.

This section includes the following topics:

Zone Implementation

All switches in the Cisco MDS 9000 Family automatically support the following basic zone features (no additional configuration is required):

- Zones are contained in a VSAN.
- Hard zoning cannot be disabled.
- Name server queries are soft-zoned.
- Only active zone sets are distributed.

- Unzoned devices cannot access each other.
- A zone or zone set with the same name can exist in each VSAN.
- Each VSAN has a full database and an active database.
- Active zone sets cannot be changed, without activating a full zone database.
- Active zone sets are preserved across switch reboots.
- Changes to the full database must be explicitly saved.
- Zone reactivation (a zone set is active and you activate another zone set) does not disrupt existing traffic.

If required, you can additionally configure the following zone features:

- Propagate full zone sets to all switches on a per VSAN basis.
- Change the default policy for unzoned members.
- Interoperate with other vendors by configuring a VSAN in the interop mode. You can also configure one VSAN in the interop mode and another VSAN in the basic mode in the same switch without disrupting each other.
- Bring E ports out of isolation.

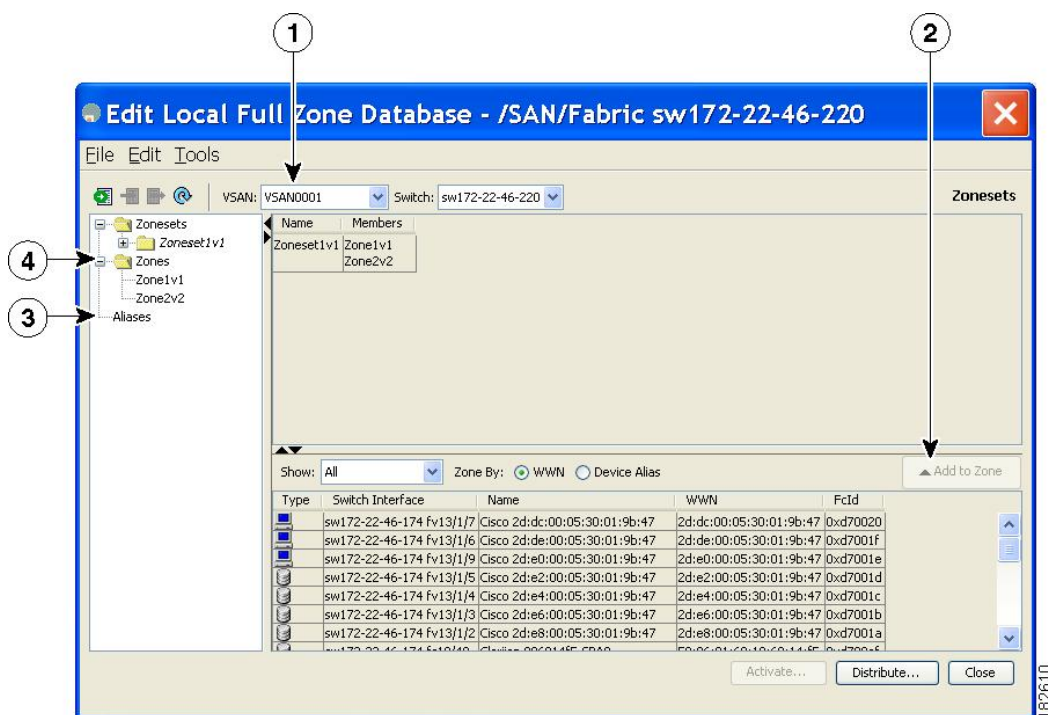
About the Edit Local Full Zone Database Tool

You can use the Edit Full Zone Database Tool to complete the following tasks:

- You can display information by VSAN by using the pull-down menu without having to get out of the screen, selecting a VSAN, and re-entering.
- You can use the **Add to zone or alias** button to move devices up or down by alias or by zone.
- You can add zoning characteristics based on alias in different folders.
- You can triple-click to rename zone sets, zones, or aliases in the tree.

The Edit Local Full Zone Database 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 43: Edit Local Full Zone Database Dialog Box, on page 406](#)).

Figure 43: Edit Local Full Zone Database Dialog Box



1	You can display information by VSAN by using the drop-down menu without closing the dialog box, selecting a VSAN, and re-entering.	3	You can add zoning characteristics based on alias in different folders.
2	You can use the Add to zone 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.

**Note**

The Device Alias radio button is visible only if device alias is in enhanced mode. For more information, see [Configuring a Zone, on page 417](#).

About Zone Sets

Zones provide a method for specifying access control. Zone sets are a grouping of zones to enforce access control in the fabric.

Zone sets are configured with the names of the member zones and the VSAN (if the zone set is in a configured VSAN).

Zone Set Distribution—You can distribute full zone sets using one of two methods: one-time distribution or full zone set distribution.

Zone Set Duplication—You can make a copy of a zone set and then edit it without altering the original 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.

About Zone Set Creation

In [Figure 44: Hierarchy of Zone Sets, Zones, and Zone Members](#), on page 407, two separate sets are created, each with its own membership hierarchy and zone members.

Figure 44: Hierarchy of Zone Sets, Zones, and Zone Members



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).

About the Default Zone

Each member of a fabric (in effect a device attached to an Nx port) can belong to any zone. If a member is not part of any active zone, it is considered to be part of the default zone. Therefore, if no zone set is active in the fabric, all devices are considered to be in the default zone. Even though a member can belong to multiple zones, a member that is part of the default zone cannot be part of any other zone. The switch determines whether a port is a member of the default zone when the attached port comes up.



Note

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

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



Note When the switch is initialized for the first time, no zones are configured and all members are considered to be part of the default zone. Members are not permitted to talk to each other.

Configure the default zone policy on each switch in the fabric. If you change the default zone policy on one switch in a fabric, be sure to change it on all the other switches in the fabric.



Note The default settings for default zone configurations can be changed.

The default zone members are explicitly listed when the default policy is configured as permit or when a zone set is active. When the default policy is configured as deny, the members of this zone are not explicitly enumerated when you issue the **show zoneset active** command view the active zone set.



Note The current default zoning policy in both the switches is deny. In the Cisco MDS 9222i Switch, the active zone set is `coco_isola_zoneset`. In the Cisco MDS 9513 Switch, there is no active zone set. However, because the default zoning policy is deny, the hidden active zone set is `d__efault__cfg` which causes zone merge to fail. The behavior is same between two Brocade switches.

You can change the default zone policy for any VSAN by choosing **VSANxx > Default Zone** from the DCNM-SAN menu tree and clicking the **Policies** tab. It is recommended that you establish connectivity among devices by assigning them to a nondefault zone.

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 or NL 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.
- IPv4 address—The IPv4 address of an attached device is in 32 bits in dotted decimal format along with an optional subnet mask. If a mask is specified, any device within the subnet becomes a member of the specified zone.
- IPv6 address—The IPv6 address of an attached device is in 128 bits in colon-(:) separated) hexadecimal format.
- Interface—Interface-based zoning is similar to port-based zoning because the switch interface is used to configure the zone. You can specify a switch interface as a zone member for both local and remote switches. To specify a remote switch, enter the remote switch WWN (sWWN) or the domain ID in the particular VSAN.



Tip The Cisco NX-OS software supports a maximum of 2048 aliases per VSAN.

Zone Enforcement

Zoning can be enforced in two ways: soft and hard. Each end device (N port or NL port) discovers other devices in the fabric by querying the name server. When a device logs in to the name server, the name server returns the list of other devices that can be accessed by the querying device. If an Nx port does not know about the FC IDs of other devices outside its zone, it cannot access those devices.

In soft zoning, zoning restrictions are applied only during interaction between the name server and the end device. If an end device somehow knows the FC ID of a device outside its zone, it can access that device.

Hard zoning is enforced by the hardware on each frame sent by an Nx port. As frames enter the switch, source-destination IDs are compared with permitted combinations to allow the frame at wirespeed. Hard zoning is applied to all forms of zoning.



Note Hard zoning enforces zoning restrictions on every frame, and prevents unauthorized access.

Switches in the Cisco MDS 9000 Family support both hard and soft zoning.

Zone Set Distribution

You can distribute full zone sets using one of two methods: one-time distribution at the EXEC mode level or full zone set distribution at the configuration mode level.

You can distribute full zone sets using one of two methods: one-time distribution or full zone set distribution.

[Table 47: Zone Set Distribution Command Differences](#), on page 409 lists the differences between these distribution methods.

Table 47: Zone Set Distribution Command Differences

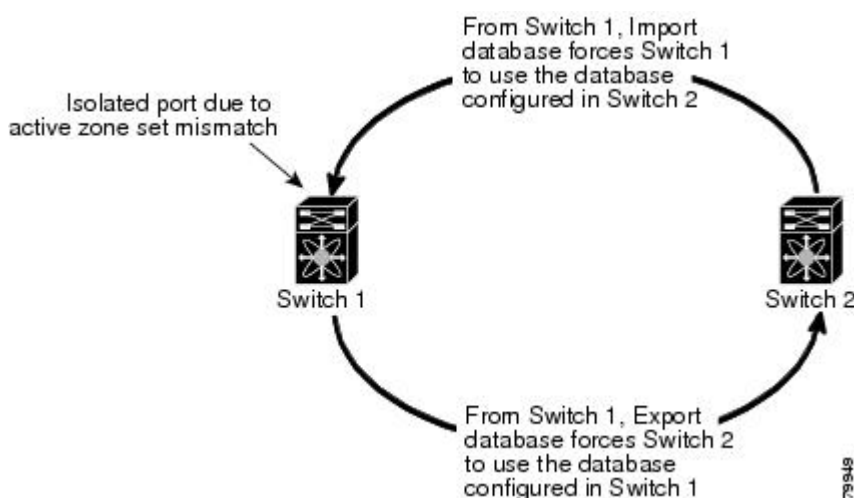
One-Time Distribution zoneset distribute vsan Command (EXEC Mode)	Full Zone Set Distribution zoneset distribute full vsan Command(Configuration Mode)
Distributes the full zone set immediately.	Does not distribute the full zone set immediately.
Does not distribute the full zone set information along with 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.

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 45: Importing and Exporting the Database](#), on page 410).
- 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 45: Importing and Exporting the Database



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.

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.

About Zone-Based Traffic Priority

The zoning feature provides an additional segregation method to prioritize select zones in a fabric and set up access control between devices. Using this feature, you can configure the quality of service (QoS) priority as

a zone attribute. You can assign the QoS traffic priority attribute to be high, medium, or low. By default, zones with no specified priority are implicitly assigned a low priority. Refer to the *Cisco MDS 9000 NX-OS Family Quality of Service Configuration Guide* for more information.

To use this feature, you need to obtain the ENTERPRISE_PKG license (refer to the *Cisco NX-OS Family Licensing Guide*) and you must enable QoS in the switch (refer to the *Cisco MDS 9000 Family NX-OS Quality of Service Configuration Guide*).

This feature allows SAN administrators to configure QoS using a familiar data flow identification paradigm. You can configure this attribute on a zone-wide basis rather than between zone members.

**Caution**

If zone-based QoS is implemented in a switch, you cannot configure the interop mode in that VSAN.

About Broadcast Zoning

**Note**

Broadcast zoning is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter.

You can configure broadcast frames in the basic zoning mode. By default, broadcast zoning is disabled and broadcast frames are sent to all Nx ports in the VSAN. When enabled, broadcast frames are only sent to Nx ports in the same zone, or zones, as the sender. Enable broadcast zoning when a host or storage device uses this feature.

[Table 48: Broadcasting Requirements, on page 411](#) identifies the rules for the delivery of broadcast frames.

Table 48: Broadcasting Requirements

Active Zoning?	Broadcast Enabled?	Frames Broadcast?	Comments
Yes	Yes	Yes	Broadcast to all Nx ports that share a broadcast zone with the source of broadcast frames.
No	Yes	Yes	Broadcast to all Nx ports.
Yes	No	No	Broadcasting is disabled.

**Tip**

If any NL port attached to an FL port shares a broadcast zone with the source of the broadcast frame, then the frames are broadcast to all devices in the loop.

**Caution**

If broadcast zoning is enabled on a switch, you cannot configure the interop mode in that VSAN.

About LUN Zoning

Logical unit number (LUN) zoning is a feature specific to switches in the Cisco MDS 9000 Family.

**Caution**

LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure the interop mode in that switch.

A storage device can have multiple LUNs behind it. If the device port is part of a zone, a member of the zone can access any LUN in the device. With LUN zoning, you can restrict access to specific LUNs associated with a device.

**Note**

When LUN 0 is not included within a zone, control traffic to LUN 0 (for example, REPORT_LUNS, INQUIRY) is supported, but data traffic to LUN 0 (for example, READ, WRITE) is denied.


- Host H1 can access LUN 2 in S1 and LUN 0 in S2. It cannot access any other LUNs in S1 or S2.
- Host H2 can access LUNs 1 and 3 in S1 and only LUN 1 in S2. It cannot access any other LUNs in S1 or S2.

**Note**

Unzoned LUNs automatically become members of the default zone.

Figure 16-4 shows a LUN-based zone example.

Figure 46: LUN Zoning Access



About Read-Only Zones

By default, an initiator has both read and write access to the target's media when they are members of the same Fibre Channel zone. The read-only zone feature allows members to have only read access to the media within a read-only Fibre Channel zone.

You can also configure LUN zones as read-only zones. Any zone can be identified as a read-only zone. By default all zones have read-write permission unless explicitly configured as a read-only zone.

About Enhanced Zoning

Table 49: Advantages of Enhanced Zoning , on page 413 lists the advantages of the enhanced zoning feature in all switches in the Cisco MDS 9000 Family.

Table 49: Advantages of Enhanced Zoning

Basic Zoning	Enhanced Zoning	Enhanced Zoning Advantages
Administrators can make simultaneous configuration changes. Upon activation, one administrator can overwrite another administrator's changes.	Performs all configurations within a single configuration session. When you begin a session, the switch locks the entire fabric to implement the change.	One configuration session for the entire fabric to ensure consistency within the fabric.
If a zone is part of multiple zone sets, you create an instance of this zone in each zone set.	References to the zone are used by the zone sets as required once you define the zone.	Reduced payload size as the zone is referenced. The size is more pronounced with bigger databases.
The default zone policy is defined per switch. To ensure smooth fabric operation, all switches in the fabric must have the same default zone setting.	Enforces and exchanges the default zone setting throughout the fabric.	Fabric-wide policy enforcement reduces troubleshooting time.
To retrieve the results of the activation on a per switch basis, the managing switch provides a combined status about the activation. It does not identify the failure switch.	Retrieves the activation results and the nature of the problem from each remote switch.	Enhanced error reporting eases the troubleshooting process.
To distribute the zoning database, you must reactivate the same zone set. The reactivation may affect hardware changes for hard zoning on the local switch and on remote switches.	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 MDS-specific zone member types (IPv4 address, IPv6 address, symbolic node name, and other types) may be used by other non-Cisco switches. During a merge, the MDS-specific types can be misunderstood by the non-Cisco switches.	Provides a vendor ID along with a vendor-specific type value to uniquely identify a member type.	Unique vendor type.
The fWWN-based zone membership is only supported in Cisco interop mode.	Supports fWWN-based membership in the standard interop mode (interop mode 1).	The fWWN-based member type is standardized.

Merging the Database

The merge behavior 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 50: Database Zone Merge Status](#), on page 414.

Table 50: Database Zone Merge Status

Local Database	Adjacent Database	Merge Status	Results of the Merge
The databases contain zone sets with the same name ³⁰ 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.

³⁰ 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.

**Caution**

Remove all non-pWWN-type zone entries on all MDS switches running Cisco SAN-OS prior to merging fabrics if there is a Cisco MDS 9020 switch running FabricWare in the adjacent fabric.

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.

Smart Zoning

Smart zoning supports zoning among more devices by reducing the number of zoning entries that needs to be programmed by considering device type information without increasing the size of the zone set. Smart zoning enables you to select the end device type. You can select if the end device type should be a host or a target. Smart zoning can be enabled at zone level, zone set level, member, and at VSAN level.



Note If smart zoning is set at the VSAN level, then you cannot enable or disable smart zoning at zone set level or zone level.

Licensing Requirements for Zoning

The following table shows the licensing requirements for this feature:

License	License Description
ENTERPRISE_PKG	The enterprise license is required to enable zoning. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

Zone Member Configuration Guidelines

All members of a zone can communicate with each other. For a zone with N members, $N * (N - 1)$ access permissions need to be enabled. Avoid configuring large numbers of targets or large numbers of initiators in a single zone. This type of configuration wastes switch resources by provisioning and managing many communicating pairs (initiator-to-initiator or target-to-target) that will never actually communicate with each other. Configuring a single initiator with a single target is the most efficient approach to zoning.

The following guidelines must be considered when creating zone members:

- Configuring only one initiator and one target for a zone provides the most efficient use of the switch resources.
- Configuring the same initiator to multiple targets is accepted.
- Configuring multiple initiators to multiple targets is not recommended.

Active and Full Zone Set Considerations

Before configuring a zone set, consider the following guidelines:

- Each VSAN can have multiple zone sets but only one zone set can be active at any given time.
- When you create a zone set, that zone set becomes a part of the full zone set.
- When you activate a zone set, a copy of the zone set from the full zone set is used to enforce zoning, and is called the active zone set. An active zone set cannot be modified. A zone that is part of an active zone set is called an active zone.
- The administrator can modify the full zone set even if a zone set with the same name is active. However, the modification will be enforced only upon reactivation.
- When the activation is done, the active zone set is automatically stored in persistent configuration. This enables the switch to preserve the active zone set information across switch resets.

- All other switches in the fabric receive the active zone set so they can enforce zoning in their respective switches.
- Hard and soft zoning are implemented using the active zone set. Modifications take effect during zone set activation.
- An FC ID or Nx port that is not part of the active zone set belongs to the default zone and the default zone information is not distributed to other switches.

**Note**

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

[Figure 47: Active and Full Zone Sets, on page 416](#) shows a zone being added to an activated zone set.

Figure 47: Active and Full Zone Sets

Read-Only Zone Configuration Guidelines

Follow these guidelines when configuring read-only zones:

- If read-only zones are implemented, the switch prevents write access to user data within the zone.
- If two members belong to a read-only zone and to a read-write zone, the read-only zone takes priority and write access is denied.
- LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure interop mode in that switch.
- Read-only volumes are not supported by some operating system and file system combinations (for example, Windows NT or Windows 2000 and NTFS file system). Volumes within read-only zones are not available to such hosts. However, if these hosts are already booted when the read-only zones are activated, then read-only volumes are available to those hosts.
- The read-only zone feature behaves as designed if either the FAT16 or FAT32 file system is used with the previously mentioned Windows operating systems.

Default Settings

[Table 51: Default Basic Zone Parameters , on page 416](#) lists the default settings for basic zone parameters.

Table 51: Default Basic Zone Parameters

Parameters	Default
Default zone policy	Denied to all members.
Full zone set distribute	The full zone set(s) is not distributed.
Zone based traffic priority	Low.
Read-only zones	Read-write attributes for all zones.

Parameters	Default
Broadcast frames	Sent to all Nx ports.
Broadcast zoning	Disabled.
Enhanced zoning	Disabled.

Configuring Zones

This section describes how to configure zones and includes the following topics:

Configuring a Zone



Tip Use the **show wwn switch** command to retrieve the sWWN. If you do not provide a sWWN, the software automatically uses the local sWWN.



Note Interface-based zoning only works with Cisco MDS 9000 Family switches. Interface-based zoning does not work if interop mode is configured in that VSAN.



Note When device aliases are used for zoning in web GUI/SAN Client, end devices must be logged into the fabric thus web GUI can configure zoning using device aliases. If end nodes are not logged in, PWWN can be used for zoning.

Configuring a Zone Using the Zone Configuration Tool

To create a zone and move it into a zone set, follow these steps:

Procedure

Step 1 Click the Zone icon in the toolbar (see [Figure 48: Zone Icon, on page 417](#)).

Figure 48: Zone Icon

You see the Select VSAN dialog box.

Step 2 Select the VSAN where you want to create a zone and click OK.

You see the 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 click **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.

Step 4

Enter a zone name.

Step 5

Check one of the following check boxes:

- a) **name="">>Read Only**—The zone permits read and denies write.
- b) **name="">>Permit QoS traffic with Priority**—You set the priority from the drop-down menu.
- c) **name="">>Restrict Broadcast Frames to Zone Members**

Step 6

Select the Smart Zoning check box to enable smart zoning.

Step 7

Click **OK** to create the zone.

If you want to move this zone into an existing zone set, skip to Step 9.

Step 8

Click **Zoneset** in the left pane and click the **Insert** icon to create a zone set.

You see the Zoneset Name dialog box.

Step 9

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 10

Select 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.

Step 11

Click **Add** to add the zone.

Adding Zone Members

Once 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, follow these steps:

Procedure

Step 1

Choose **Zone > Edit Local Full Zone Database**.

You see the Select VSAN dialog box.

Step 2

Select a VSAN and click **OK**.

You see the Edit Local Full Zone Database dialog box for the selected VSAN.

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.

Note The Device Alias radio button is visible only if device alias is in enhanced mode. For more information, see [Configuring a Zone, on page 417](#).

- Step 4** Click the browse button and select a port name or check the **LUN** check box and click the browse button to configure LUNs.
- Step 5** Select the options for Device Type field. You can select any one of the options: Host, Storage, or Both.
- Step 6** 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.
-

Filtering End Devices Based on Name, WWN, or FC ID

To filter the end devices and device aliases, follow these steps:

Procedure

- Step 1** Click the Zone icon in the toolbar.
- Step 2** Select Name, WWN, or FC ID from the With drop-down list.
- Step 3** Enter a filter condition, such as *zo1*, in the Filter text box.
- Step 4** Click **Go**.
-

Adding Multiple End Devices to Multiple Zones

To add multiple end devices to multiple zones, follow these steps:

Procedure

- Step 1** Click the Zone icon in the toolbar.
- Step 2** Press the Control key to select multiple end devices.
- Step 3** Right-click the device and then select **Add to Zone**.
- Step 4** Press the Control key to select multiple zones from the pop-up window displayed.
- Step 5** Click **Add**.
- Selected end devices are added to the selected zones.
-

Using the Quick Config Wizard



Note The Quick Config Wizard supports only switch interface zone members.

As of Cisco SAN-OS Release 3.1(1) and NX-OS Release 4.1(2), you can use the Quick Config Wizard on the Cisco MDS 9124 Switch to add or remove zone members per VSAN. You can use the Quick Config

Wizard to perform interface-based zoning and to assign zone members for multiple VSANs using Device Manager.



Note The Quick Config Wizard is supported on Cisco MDS 9124, MDS 9134, MDS 9132T, MDS 9148, MDS 9148S, MDS 9148T, MDS 9396S, and MDS 9396T fabric switches, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

The Quick Config Wizard can only be used on standalone switches that do not have any existing zoning defined on the switch.

To add or remove ports from a zone and to zone only the devices within a specific VSAN using Device Manager on the Cisco MDS 9124 Switch, follow these steps:

Procedure

-
- Step 1** Choose **FC > Quick Config** or click the Zone icon in the toolbar.
- You see the Quick Config Wizard (see [Figure 50: Quick Config Wizard, on page 420](#)) with all controls disabled and the Discrepancies dialog box (see [Figure 49: Discrepancies Dialog Box, on page 420](#)), which shows all unsupported configurations.
- Note** You will see the Discrepancies dialog box only if there are any discrepancies.
- Figure 49: Discrepancies Dialog Box*
- Step 2** Click **OK** to continue.
- You see the Quick Config Wizard dialog box (see [Figure 50: Quick Config Wizard, on page 420](#)).
- Caution** If there are discrepancies and you click **OK**, the affected VSANs in the zone databases are cleared. This might be disruptive if the switch is in use.
- Figure 50: Quick Config Wizard*
- Step 3** Check the check box in the **Ports Zoned To** column for the port you want to add or remove from a zone. The check box for the matching port is similarly set. The selected port pair is added or removed from the zone, which creates a two-device zone.
- The VSAN drop-down menu provides a filter that enables you to zone only those devices within a selected VSAN.
- Step 4** Right-click any of the column names to show or hide a column.
- Step 5** Click **Next** to verify the changes.
- You see the Confirm Changes dialog box.
- Step 6** If you want to see the CLI commands, right-click in the dialog box and click **CLI Commands** from the pop-up menu.
- Step 7** Click **Finish** to save the configuration changes.
-

Configuring Zone Sets

This section describes how to configure zones and includes the following topics:

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, follow these steps:

Procedure

-
- | | |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Choose Zone > Edit Local Full Zone Database .
You see the Select VSAN dialog box. |
| Step 2 | Select 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 pre-activation check dialog box. |
| Step 4 | Click Yes to review the differences.
You see the Local vs. Active Differences dialog box. |
| Step 5 | Click Close to close the dialog box.
You see the 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.
You see the Zone Log dialog box, which shows if the zone set activation was successful. |
-

Deactivating a Zone Set

To deactivate an existing zone set, follow these steps:

Procedure

-
- | | |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Right-click the zone set you want to deactivate, and then click Deactivate from the pop-up menu.
You see the Deactivate Zoneset dialog box. |
| Step 2 | Enter deactivate in the text box, and then click OK.
You see the Input dialog box. |
| Step 3 | Enter deactivate in the text box, and then click OK to deactivate the zone set. |

Note To enable this option, you need to modify the `server.properties` file. Refer to the *Cisco DCNM Fundamentals Guide* to know more about modifying `server.properties` file.

Displaying Zone Membership Information

To display zone membership information for members assigned to zones, follow these steps:

Procedure

Step 1 Choose **Zone > Edit Local Full Zone Database**.

You see the Select VSAN dialog box.

Step 2 Select 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.

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.

Configuring the Default Zone Access Permission

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

Procedure

Step 1 Expand a **VSAN** and then select **Default Zone** in the DCNM-SAN Logical Domains pane.

Step 2 Click the **Policies** tab in the Information pane.

You see the zone policies information in the Information pane.

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 Behaviour field, choose either **permit** or **deny** from the drop-down menu.

Creating FC Aliases

To create an FC alias, follow these steps:

Procedure

- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
- Step 3** Click **Aliases** in the lower left pane. The right pane lists the existing aliases.
- Step 4** Click the **Insert** icon to create an alias.
You see the 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, follow these steps:

Procedure

- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
- Step 3** Select the member(s) you want to add from the Fabric pane and click **Add to Alias** or click the alias where you want to add members and click the **Insert** icon.
You see the Add Member to Alias dialog box.
- Note** The Device Alias radio button is visible only if device alias is in enhanced mode. For more information, see [Configuring a Zone, on page 417](#).
- Step 4** Click the browse button and select a port name or check the **LUN** check box and click the browse button to configure LUNs.
- Step 5** 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 card or switch is changed in your fabric.

To convert switch port and FC ID members to pWWN members, follow these steps:

Procedure

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Select 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, which lists 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.
-

Creating Zone Sets and Adding Member Zones

The pWWN of the virtual target does not appear in the zoning end devices database in DCNM-SAN. 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 DCNM-SAN 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, on page 418](#).

Set the device alias mode to **enhanced** when using SDV (because the pWWN of a virtual device could change). For example, SDV is enabled on a switch and a virtual device is defined. SDV assigns a pWWN for the virtual device, and it is zoned based on the pWWN in a zone. If you later disable SDV, this configuration is lost. If you reenables SDV and create the virtual device using the same name, there is no guarantee that it will get the same pWWN again. You will have to rezone the pWWN-based zone. However, if you perform zoning based on the device-alias name, there are no configuration changes required if or when the pWWN changes. Be sure you understand how device alias modes work before enabling them. Refer to [Chapter 26, “Distributing Device Alias Services.”](#) for details and requirements about device alias modes.



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



Tip You do not have to copy the running configuration to the startup configuration. Issue the **copy running-config startup-config** command to store the active zone set. However, you need to copy the running configuration to the startup configuration issue the **copy running-config startup-config** command to explicitly store full zone sets. It is not available across switch resets.

**Caution**

If you deactivate the active zone set in a VSAN that is also configured for IVR, the active IVR zone set (IVZS) is also deactivated and all IVR traffic to and from the switch is stopped. This deactivation can disrupt traffic in more than one VSAN. Before deactivating the active zone set, check the active zone analysis for the VSAN (see the [“Zone and Zone Set Analysis” section on page 16-61](#)). To reactivate the IVZS, you must reactivate the regular zone set (refer to the *Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide*).

**Caution**

If the currently active zone set contains IVR zones, activating the zone set from a switch where IVR is not enabled disrupts IVR traffic to and from that VSAN. We strongly recommend that you always activate the zone set from an IVR-enabled switch to avoid disrupting IVR traffic.

Filtering Zones, Zone Sets, and Device Aliases Based on Name

To filter the zones, zone sets, or device aliases, follow these steps:

Procedure

- Step 1** Click the Zone icon in the toolbar.
- Step 2** Enter a filter condition, such as *zo1*, in the Filter text box.
- Step 3** Click Go.

Adding Multiple Zones to Multiple Zone Sets

To add multiple zones to multiple zone sets, follow these steps:

Procedure

- Step 1** Click the Zone icon in the toolbar.
 - Step 2** From the tree view, select **Zoneset**.
 - Step 3** Press the control key to select multiple zones.
 - Step 4** Right-click and then select **Add to Zoneset**.
 - Step 5** Press the control key to select multiple zone sets from the pop-up window displayed.
 - Step 6** Click **Add**.
- Selected zones are added to the selected zone sets.

Enabling Full Zone Set Distribution

All switches in the Cisco MDS 9000 Family distribute active zone sets when new E port links come up or when a new zone set is activated in a VSAN. The zone set distribution takes effect while sending merge requests to the adjacent switch or while activating a zone set.

To enable full zone set and active zone set distribution to all switches on a per-VSAN basis, follow these steps:

Procedure

-
- Step 1** Expand a **VSAN** and select a zone set in the Logical Domains pane.
You see the zone set configuration in the Information pane. The Active Zones tab is the default.
 - Step 2** Click the **Policies** tab.
You see the configured policies for the zone.
 - Step 3** In the **Propagation** column, choose full Zoneset from the drop-down menu.
 - Step 4** Click **Apply Changes** to propagate the full zone set.
-

Enabling a One-Time Distribution

Use the **zoneset distribute vsan** *vsan-id* command in EXEC mode to perform this distribution.

```
switch# zoneset distribute vsan 2
Zoneset distribution initiated. check zone status
```

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, follow these steps:

Procedure

-
- 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.
-

What to do next

This procedure command only distributes the full zone set information; it does not save the information to the startup configuration. You must explicitly save the running configuration to the startup configuration **copy running-config startup-config** command to save the full zone set information to the startup configuration.



Note The **zoneset distribute vsan vsan-id** command one-time distribution of the full zone set is supported in **interop 2** and **interop 3** modes, not in **interop 1** mode.

Importing and Exporting Zone Sets

To import or export the zone set information from or to an adjacent switch, follow these steps:

Procedure

- Step 1** Choose **Tools > Merge Fail Recovery**.
You see the Zone Merge Failure Recovery dialog box.
- Step 2** Click the **Import Active Zoneset** or the **Export Active Zoneset** radio button.
- Step 3** Select the switch from which to import or export the zone set information from the drop-down list.
- Step 4** Select the VSAN from which to import or export the zone set information from the drop-down list.
- Step 5** Select the interface to use for the import process.
- Step 6** Click **OK** to import or export the active zone set.

What to do next



Note Issue the **import** and **export** commands from a single switch. Importing from one switch and exporting from another switch can lead to isolation again.

Copying Zone Sets

On the Cisco MDS Family switches, you cannot edit an active zone set. However, you can copy an active zone set to create a new zone set that you can edit.

To make a copy of a zone set, follow these steps:

Procedure

- Step 1** Choose **Zone > Copy Full Zone Database**.
You see the 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** Select the source VSAN from the drop-down list.
- Step 4** If you selected **Copy Full**, select the source switch and the destination VSAN from those drop-down lists.
- Step 5** Select the destination switch from the drop-down list.

Step 6 Click **Copy** to copy the database.

What to do next



Caution

If the Inter-VSAN Routing (IVR) feature is enabled and if IVR zones exist in the active zone set, then a zone set copy operation copies all the IVR zones to the full zone database. To prevent copying to the IVR zones, you must explicitly remove them from the full zone set database before performing the copy operation. Refer to the *Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide* for more information on the IVR feature.

Backing Up Zones

To back up the full zone configuration, follow these steps:

Procedure

Step 1 Choose **Zone > Edit Local Full Zone Database**.

You see the Select VSAN dialog box.

Step 2 Select a VSAN and click **OK**.

You see the Edit Local Full Zone Database dialog box for the selected VSAN.

Step 3 Choose **File > Backup > This VSAN Zones** to back up the existing zone configuration to a workstation using TFTP, SFTP, SCP, or FTP.

You see the Backup Zone Configuration dialog box.

You can edit this configuration before backing up the data to a remote server.

Step 4 Provide the following Remote Options information to back up data onto a remote server:

- a) **Using**—Select the protocol.
- b) **Server IP Address**—Enter the IP address of the server.
- c) **UserName**—Enter the name of the user.
- d) **Password**—Enter the password for the user.
- e) **File Name(Root Path)**—Enter the path and the filename.

Step 5 Click **Backup** or click **Cancel** to close the dialog box without backing up.

Restoring Zones

To restore the full zone configuration, follow these steps:

Procedure

- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
- Step 3** Choose **File > Restore** to restore a saved zone configuration using TFTP, SFTP, SCP, or FTP.
You see the Restore Zone Configuration dialog box.
You can edit this configuration before restoring it to the switch.
- Step 4** Provide the following Remote Options information to restore data from a remote server:
- a) Using—Select the protocol.
 - b) Server IP Address—Enter the IP address of the server.
 - c) UserName—Enter the name of the user.
 - d) Password—Enter the password for the user.
 - e) File Name—Enter the path and the filename.
- Step 5** Click **Restore** to continue or click **Cancel** to close the dialog box without restoring.
-

What to do next



Note Click View Config to see information on how the zone configuration file from a remote server will be restored. When you click Yes in this dialog box, you are provided with the CLI commands that are executed. To close the dialog box, click Close.



Note Backup and Restore options are available to switches that run Cisco NX-OS Release 4.1(3) or later.

Renaming Zones, Zone Sets, and Aliases

To rename a zone, zone set, or alias, follow these steps:

Procedure

- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.

- 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, FC alias, or zone attribute group, follow these steps:

Procedure

- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
 - Step 2** Select 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. The default name is the word **Clone** followed by the original name.
 - 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, follow these steps:

Procedure

- 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, use the following command:


```
switch# clear zone database vsan 2
```

To clear the zone server database, refer to the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.



Note After issuing a **clear zone database** command, you must explicitly issue the **copy running-config startup-config** to ensure that the running configuration is used when the switch reboots.



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



Note After clearing the zone server database, you must explicitly **copy the running configuration to the startup configuration** to ensure that the running configuration is used when the switch reboots.

Configuring Zone-Based Traffic Priority

To configure the zone priority, follow these steps:

Procedure

- Step 1** Expand a **VSAN** and then select a zone set in the Logical Domains pane.
- Step 2** Click the **Policies** tab in the Information pane.
You see the Zone policy information in the Information pane.
- Step 3** Use the check boxes and drop-down menus to configure QoS on the default zone.
- Step 4** Click **Apply Changes** to save the changes.

Configuring Default Zone QoS Priority Attributes

QoS priority attribute configuration changes take effect when you activate the zone set of the associated zone.



Note If a member is part of two zones with two different QoS priority attributes, the higher QoS value is implemented. This situation does not arise in the VSAN-based QoS as the first matching entry is implemented.

To configure the QoS priority attributes for a default zone, follow these steps:

Procedure

- Step 1** Choose **Zone > Edit Local Full Zone Database**.

You see the Select VSAN dialog box.

Step 2 Select 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.

Step 4 Check the **Permit QoS Traffic with Priority** check box and set the QoS Priority drop-down menu to low, **medium**, or **high**.

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

Configuring the Default Zone Policy

To permit or deny traffic in the default zone, follow these steps:

Procedure

Step 1 Choose **Zone > Edit Local Full Zone Database**.

You see the Select VSAN dialog box.

Step 2 Select 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.

Step 4 Set the Policy drop-down menu to **permit** to permit traffic in the default zone, or set it to **deny** to block traffic in the default zone.

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

Configuring Smart Zoning

To configure smart zoning, follow these steps:

Procedure

Step 1 Expand a **VSAN** and then select a zone set in the Logical Domains pane.

Step 2 Click the **Smart Zoning** tab in the Information pane.

You see the smart zoning information in the Information pane.

Step 3 You can view the details under the Switch, Status, Command, Last Command, and Result headings.

Step 4 You can set the Status, and Command fields.

Step 5 Click **Apply Changes** to save these changes.

Configuring Global Zone Policies

To configure global zone policy, follow these steps:

Procedure

- Step 1** In the Logical Domains pane, select ALL VSANs.
 - Step 2** Click the **Global Zone Policies** tab in the Information pane.
You see the Global Zone Policy information in the Information pane.
 - Step 3** Set the type of switch under the Switch column.
 - Step 4** You either Deny or Permit the Zone Behaviour and set the Propagation Mode.
 - Step 5** Select if the Smart Zoning feature is enabled or disabled.
 - Step 6** Click **Apply Changes** to save these changes.
-

Configuring Broadcast Zoning

To broadcast frames in the basic zoning mode, follow these steps:

Procedure

- Step 1** Expand a **VSAN** and then select a zone set in the Logical Domains pane.
 - Step 2** Click the **Policies** tab in the Information pane.
You see the Zone policy information in the Information pane.
 - Step 3** Check the **Broadcast** check box to enable broadcast frames on the default zone.
 - Step 4** Click **Apply Changes** to save these changes.
-

Configuring a LUN-Based Zone

To configure a LUN-based zone, follow these steps:

Procedure

- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
- Step 3** Click the zone where you want to add members and click the **Insert** icon.
You see the Add Member to Zone dialog box.

- Step 4** Click either the **WWN** or **FCID** radio button from the Zone By options to create a LUN-based zone.
- Step 5** Check the **LUN** check box and click the browse button to configure LUNs.
- Step 6** Click **Add** to add this LUN-based zone.

Assigning LUNs to Storage Subsystems

LUN masking and mapping restricts server access to specific LUNs. If LUN masking is enabled on a storage subsystem and if you want to perform additional LUN zoning in a Cisco MDS 9000 Family switch, obtain the LUN number for each host bus adapter (HBA) from the storage subsystem and then configure the LUN-based zone procedure provided in the [Configuring a LUN-Based Zone, on page 433](#).



Note

Refer to the relevant user manuals to obtain the LUN number for each HBA.



Caution

If you make any errors when assigning LUNs, you might lose data.

Configuring Read-Only Zones

To configure read-only zones, follow these steps:

Procedure

- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Select 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 and click the **Insert** icon to add a zone.
You see the Create Zone Dialog Box.
- Step 4** Check the **Read Only** check box to create a read-only zone.
- Step 5** Click **OK**.

What to do next



Note

To configure the **read-only** option for a default zone, see [Configuring the Default Zone Policy, on page 432](#).

Changing from Basic Zoning to Enhanced Zoning

To change to the enhanced zoning mode from the basic mode, follow these steps:

Procedure

- 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 be able to automatically start a session, acquire a fabric wide lock, distribute the active and full zoning database using the enhanced zoning data structures, distribute zoning policies, and then release the lock. All switches in the fabric then move to the enhanced zoning mode.
- Tip** After moving from basic zoning to enhanced zoning, we recommend that you save the running configuration.
-

Changing from Enhanced Zoning to Basic Zoning

The standards do not allow you to move back to basic zoning. However, Cisco MDS switches allow this move to enable you to downgrade and upgrade to other Cisco SAN-OS or Cisco NX-OS releases.

To change to the basic zoning mode from the enhanced mode, follow these steps:

Procedure

- Step 1** Verify that the active and full zone set do not contain any configuration that is specific to the enhanced zoning mode.
- If such configurations exist, delete them before proceeding with this procedure. If you do not delete the existing configuration, the Cisco NX-OS software automatically removes them.
- Step 2** Set the operation mode to basic zoning mode.
- You will be able to automatically start a session, acquire a fabric wide lock, distribute the zoning information using the basic zoning data structure, apply the configuration changes, and release the lock from all switches in the fabric. All switches in the fabric then move to basic zoning mode.
- Note** If a switch running Cisco SAN-OS Release 2.0(1b) and NX-OS 4(1b) or later, with enhanced zoning enabled is downgraded to Cisco SAN-OS Release 1.3(4), or earlier, the switch comes up in basic zoning mode and cannot join the fabric because all the other switches in the fabric are still in enhanced zoning mode.
-

Enabling Enhanced Zoning

By default, the enhanced zoning feature is disabled in all switches in the Cisco MDS 9000 Family.

To enable enhanced zoning in a VSAN, follow these steps:

Procedure

- Step 1** Expand a VSAN and then select a zone set in the Logical Domains pane.
You see the zone set configuration in the Information pane.
 - Step 2** Click the **Enhanced** tab.
You see the current enhanced zoning configuration.
 - Step 3** From the Action drop-down menu, choose **enhanced** to enable enhanced zoning in this VSAN.
 - Step 4** Click **Apply Changes** to save these changes.
-

Modifying the Zone Database

Modifications to the zone database is done within a session. A session is created at the time of the first successful configuration command. On creation of a session, a copy of the zone database is created. Any changes done within the session are performed on this copy of the zoning database. These changes in the copy zoning database are not applied to the effective zoning database until you commit the changes. Once you apply the changes, the session is closed.

If the fabric is locked by another user and for some reason the lock is not cleared, you can force the operation and close the session. You must have permission (role) to clear the lock in this switch and perform the operation on the switch from where the session was originally created.

Analyzing a Zone Merge

To perform a zone merge analysis, follow these steps:

Procedure

- Step 1** Choose **Zone > Merge Analysis**.
You see the Zone Merge Analysis dialog box.
 - Step 2** Select the first switch to be analyzed from the Check Switch 1 drop-down list.
 - Step 3** Select the second switch to be analyzed from the And Switch 2 drop-down list.
 - Step 4** Enter the VSAN ID where the zone set merge failure occurred in the For Active Zoneset Merge Problems in VSAN Id field.
 - Step 5** Click **Analyze** to analyze the zone merge.
 - Step 6** Click **Clear** to clear the analysis data in the Zone Merge Analysis dialog box.
-

Preventing Zones From Flooding FC2 Buffers

By using the **zone fc2 merge throttle enable** command you can throttle the merge requests that are sent from zones to FC2 and prevent zones from flooding FC2 buffers. This command is enabled by default. This command

can be used to prevent any zone merge scalability problem when you have a lot of zones. Use the **show zone status** command to view zone merge throttle information.

Broadcasting a Zone

You can specify an enhanced zone to restrict broadcast frames generated by a member in this zone to members within that zone. Use this feature when the host or storage devices support broadcasting.

Table 52: [Broadcasting Requirements](#) , on page 437 identifies the rules for the delivery of broadcast frames.

Table 52: Broadcasting Requirements

Active Zoning?	Broadcast Enabled?	Frames Broadcast?	Comments
Yes	Yes	Yes	Broadcast to all Nx ports that share a broadcast zone with the source of broadcast frames.
No	Yes	Yes	Broadcast to all Nx ports.
Yes	No	No	Broadcasting is disabled.



Tip If any NL port attached to an FL port shares a broadcast zone with the source of the broadcast frame, then the frames are broadcast to all devices in the loop.

Configuring System Default Zoning Settings

You can configure default settings for default zone policies, full zone distribution, and generic service permissions for new VSANs on the switch.



Note Since VSAN 1 is the default VSAN and is always present on the switch, the **system default zone** commands have no effect on VSAN 1.

Configuring Zone Generic Service Permission Settings

Zone generic service permission setting is used to control zoning operation through generic service (GS) interface. The zone generic service permission can be read-only, read-write or none (deny).

Compacting the Zone Database for Downgrading

Prior to Cisco SAN-OS Release 3.0(1), only 2000 zones are supported per VSAN. If you add more than 2000 zones to a VSAN, a configuration check is registered to indicate that downgrading to a previous release could cause you to lose the zones over the limit. To avoid the configuration check, delete the excess zones and compact the zone database for the VSAN. If there are 2000 zones or fewer after deleting the excess zones, the compacting process assigns new internal zone IDs and the configuration can be supported by Cisco SAN-OS Release 2.x or earlier. Perform this procedure for every VSAN on the switch with more than 2000 zones.



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, refer to the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.

Displaying Zone Information

You can view any zone information by using the **show** command. If you request information for a specific object (for example, a specific zone, zone set, VSAN, or alias, or keywords such as **brief** or **active**), only information for the specified object is displayed. If you do not request specific information, all available information is displayed.

To view zone information and statistics, follow these steps:

Procedure

- Step 1** Expand a **VSAN** and select a zone set in the Logical Domains pane.
You see the zone configuration in the Information pane.
- Step 2** Click the **Read Only Violations, Statistics** tab or the **LUN Zoning Statistics** tab to view statistics for the selected zone.

Configuration Examples for Zoning

[Figure 51: Fabric with Two Zones](#), on page 438 illustrates a zone set with two zones, zone 1 and zone 2, in a fabric. Zone 1 provides access from all three hosts (H1, H2, H3) to the data residing on storage systems S1 and S2. Zone 2 restricts the data on S3 to access only by H3. Note that H3 resides in both zones.

Figure 51: Fabric with Two Zones

You can partition this fabric into zones using other methods. [Figure 52: Fabric with Three Zones](#), on page 438 illustrates another possibility. Assume that there is a need to isolate storage system S2 for the purpose of testing new software. To achieve this, zone 3 is configured, which contains only host H2 and storage S2. You can restrict access to just H2 and S2 in zone 3, and to H1 and S1 in zone 1.

Figure 52: Fabric with Three Zones

Field Descriptions for Zones

The following are the field descriptions for zoning.

Zone Set Active Zones

Field	Description
Zone	Zone name.
Type	Zone member type.
Device Type	Specifies if the end device type is host, storage, or both.
Switch Interface	Switch interface to which the zone member is connected to.
Name	Zone member name.
WWN	Zone member WWN.
FcId	Zone member FC ID.
LUNs	Zone member LUN.
Status	<ul style="list-style-type: none"> • Not in Fabric: If zone member is not in the fabric. • Not in VSAN: If zone member is not present in the VSAN. • n/a: Cannot determine status. • Empty: Member is present in fabric and correct VSAN and can communicate with other members of the zone.

Zone Set Unzoned

Field	Description
Name	Zone member name.
WWN	Zone member WWN.
FcId	Zone member FC ID.

Zone Set Status

Field	Description
Status	Indicates the outcome of the most recent activation or deactivation.
Activation Time	When this entry was most recently activated. If this entry has been activated prior to the last reinitialization of the local network management system, then this value will be N/A.
FailureCause	The reason for the failure of the zone set activation or deactivation.
FailedSwitch	The domain ID of the device in the fabric that has caused the Change Protocol to fail.

Field	Description
Active == Local?	Indicates whether the enforced database is the same as the local database on this VSAN. If true, then they are the same. If false, then they are not the same.
Active Zoneset	The name of the enforced IV zone set.
Hard Zoning	Indicates whether the hard zoning is enabled on this VSAN. Hard zoning is a mechanism by which zoning is enforced in hardware. If true, then hard zoning is enabled on this VSAN. If false, then hard zoning is not enabled on this VSAN.

Zone Set Policies

Field	Description
Default Zone Behavior	Controls the behavior of the default zone on this VSAN. If it is set to permit, then the members of the default zone on this VSAN can communicate with each other. If it is set to deny, then the members of the default zone on this VSAN cannot communicate with each other.
Default Zone ReadOnly	Indicates whether SCSI read operations are allowed on members of the default zone which are SCSI targets, on this VSAN. If true, then only SCSI read operations are permitted. So, this default zone becomes a read-only default zone on this VSAN. If false, then both SCSI read and write operations are permitted.
Default Zone QoS	Specifies whether the QoS attribute for the default zone on this VSAN is enabled. If true, then QoS attribute for the default zone on this VSAN is enabled. If false, then the QoS attribute for the default zone on this VSAN is disabled.
Default Zone QoS Priority	Specifies the QoS priority value.
Default Zone Broadcast	Specifies if broadcast zoning is enabled on this default zone on this VSAN. If true, then it is enabled. If false, then it is disabled.
Smart Zoning	Specifies if the smart zoning feature is enabled or disabled at the VSAN level
Propagation	Controls the way zoneset information is propagated during Merge/Change protocols on this VSAN
Read From	Specifies whether the management station wishes to read from the effective database or from the copy database.

Zone Set Active Zones Attributes

Field	Description
Name	Zone name.
Read Only	Indicates if only SCSI read operations are allowed on members of the default zone which are SCSI targets on this VSAN. If true, then only SCSI read operations are permitted. So, this default zone becomes a read-only default zone on this VSAN. If false, then both SCSI read and write operations are permitted.

Field	Description
QoS	Specifies whether the QoS attribute for the default zone on this VSAN is enabled. If true, then QoS attribute for the default zone on this VSAN is enabled. If false, then the QoS attribute for the default zone on this VSAN is disabled.
QoS Priority	Specifies QoS priority value (Low, Medium, or High).
Broadcast	Specifies if broadcast zoning is enabled on this default zone on this VSAN. If true, then it is enabled. If false, then it is disabled.
Smart Zoning	Specifies if the smart zoning feature is enabled. on this VSAN. If the check box is unchecked, then it is disabled.

Zone Set Enhanced

Field	Description
Action	When set to basic(1), results in the zone server operating in the basic mode as defined by FC-GS4 standards. When set to enhanced(2), results in the zone server operating in the enhanced mode as defined by FC-GS4 standards.
Result	The outcome of setting the mode of operation of the local zone server on this VSAN.
Config DB Locked By	Specifies the owner for this session.
Config DB Discard Changes	Assists in committing or clearing the contents of the copy database on this session.
Config DB Result	Indicates the outcome of setting the corresponding instance of czseSessionCntl to commitChanges(1).
Enforce Full DB Merge	Controls the zone merge behavior. If this object is set to allow, then the merge takes place according to the merge rules. If set to restrict, then if the merging databases are not exactly identical, the Inter-Switch Link (ISL) between the devices is isolated.
Read From	Specifies whether the management station wishes to read from the effective database or from the copy database.

Smart Zoning

Field	Description
Switch	Specifies the type of which where smart zoning feature exists.
Status	Specifies if the smart zoning feature is enabled or disabled.
Command	Specifies the switch level command for smart zoning. If the command is disabled in one switch then smart zoning will be disabled in the whole fabric.

Field	Description
Last Command	Specified the previous command mode of the switch. Enabled or disabled.
Result	Specifies if the enable or disable action has been successful or unsuccessful.

Zone Set Read Only Violations

Field	Description
Violations	The number of data-protected Check Condition error responses sent by the local zone server.

Zone Set Statistics

Field	Description
Merge Req Tx	The number of merge request frames sent by this zone server to other zone servers in the fabric on this VSAN.
Merge Req Rx	The number of merge request frames received by this zone server from other zone servers in the fabric on this VSAN.
Merge Acc Tx	The number of merge accept frames sent by this zone server to other zone servers in the fabric on this VSAN.
Merge Acc Rx	The number of merge accept frames received by this zone server from other zone servers in the fabric on this VSAN.
Change Req Tx	The number of change requests sent by this zone server to other zone servers in the fabric on this VSAN.
Change Req Rx	The number of change requests received by this zone server from other zone servers in the fabric on this VSAN.
Change Acc Tx	The number of change responses sent by this zone server to other zone servers in the fabric on this VSAN.
Change Acc Rx	The number of change responses received by this zone server from other zone servers in the fabric on this VSAN.
GS3 Rej Tx	The number of GS3 requests rejected by this zone server on this VSAN.
GS3 Req Rx	The number of GS3 requests received by this zone server on this VSAN.

Zone Set LUN Zoning Statistics

Field	Description
INQUIRY	The number of SCSI INQUIRY commands that have been received by the local zone server.

Field	Description
REPORT LUN	The number of SCSI Report LUNs commands that have been received by the local zone server. Typically the Report LUNs command is sent only for LUN 0.
SENSE	The number of SCSI SENSE commands that have been received by the local zone server.
Other Cmds	The number of SCSI Read, Write, Seek commands received by the local zone server.
BadInquiry Errors	The number of No LU error responses sent by the local zone server.
Illegal Errors	The number of Illegal Request Check Condition responses sent by the local zone server.

Zone Set Members

Field	Description
Zone	Default zone.
Type	FCID.
Device Type	Specifies if the end device type is host, storage, or both.
Switch Interface	Switch interface to which the zone member is connected to.
Name	Zone member name.
WWN	Zone member WWN.
FcId	Zone member FC ID.
Luns	Zone member LUN.
Status	<ul style="list-style-type: none"> • Not in Fabric: If zone member is not in the fabric. • Not in VSAN: If zone member is not present in the VSAN. • n/a: Cannot determine status. <p>Empty: Member is present in fabric and correct VSAN and can communicate with other members of the zone.</p>



CHAPTER 17

Configuring FCoE

- [Configuring FCoE, on page 445](#)

Configuring FCoE

This chapter describes how to configure Fibre Channel over Ethernet (FCoE) on a Cisco Nexus 5000 Series Switch, Cisco Nexus 7000 Series Switch, and Cisco 9000 Family MDS switch.

About FCoE

Cisco Nexus 5000 Series Switch, Cisco Nexus 7000 Series Switch, and Cisco MDS 9000 family switches support Fibre Channel over Ethernet (FCoE), which allows Fibre Channel and Ethernet traffic to be carried on the same physical Ethernet connection between the switch and the servers. FCoE requires the underlying Ethernet to be full duplex and to provide lossless behavior for Fibre Channel traffic.

The FCoE Initialization Protocol (FIP) allows the switch to discover and initialize FCoE-capable entities that are connected to an Ethernet LAN.

Guidelines and Limitations

When configuring FCoE, note the following guidelines and limitations:

- FCoE is supported on 10-Gigabit Ethernet interfaces.
- FCoE is not supported on private VLANs.
- DPVM supports MAC-based device mapping for FCoE devices. DPVM does not support pWWN mapping for FCoE devices.

Configuring FCoE

Enabling FCoE

Fibre Channel over Ethernet (FCoE) provides a method of transporting Fibre Channel traffic over a physical Ethernet connection. By default, each Ethernet interface attempts to enable FCoE by advertising that it has

FCoE to the adapter. If the FCoE negotiation fails, you can configure the Cisco Nexus 5000 Series switch to disable FCoE for this interface.



Note In Cisco Nexus 5000 Series switches, FCoE is supported on all 10-Gigabit Ethernet interfaces.

To enable or disable FCoE features on a switch using Device Manager, follow these steps:

Procedure

Step 1 Launch Device Manager from the Cisco Nexus 5000 Series switch.

Note Use the Control tab to enable FCoE on a Cisco Nexus 5000 Series switch.

Step 2 Choose **Admin > Feature Control**.

You see the Feature Control dialog box.

Note You cannot enable FCoE using Device Manager on Cisco Nexus 7000 series and Cisco MDS 900 family switches. Cisco Nexus 7000 series and Cisco MDS 9000 Family switches uses a feature set to display FCoE information.

Step 3 In the dialog box, in the table, click the **fcoe_mgr** row, and then click the **Action** cell in the fcoe_mgr row. From the drop-down list, choose **enable** to enable the FCoE feature in the switch.

Note You can also disable the FCoE feature in the switch. To do so, from the drop-down list in the Action column, choose **disable**.

Step 4 Click **Apply**.

Note If the Cisco Nexus 5000 Series switch is running a Cisco NX-OS release prior to Release 4.2(1), you must do the following after you enable or disable FCoE on the switch:

- a) In the confirmation dialog box that appears, click **Yes** to enable the FCoE feature in the switch.
- b) Reboot the switch before you use the FCoE feature.

Configuring FCoE Using DCNM for SAN

From Cisco NX-OS Release 5.2, FCoE is supported on MDS and Cisco Nexus 7000 switches. To enable or disable FCoE, Cisco MDS 9000 Family and Cisco Nexus 7000 switches uses feature set MIBs.

To configure FCoE on a switch, follow these steps:

Procedure

Step 1 In the Physical Attributes pane, choose **Switches > FC Services > FCoE**.

The Config tab displays the FCoE parameters for each Cisco Nexus 5000 Series, Cisco Nexus 7000 Series, and Cisco MDS 9000 Family switches. Use the VLAN-VSAN mapping tab to create mappings. [Table 53: FCoE Parameters](#), on page 447 lists the FCoE parameters for a switch.

For more information on configuring Cisco Nexus 5000 Series and Nexus 7000 Series switches, see the Cisco Nexus 5000 Series and Nexus 7000 Series Configuration Guides.

Table 53: FCoE Parameters

Parameter	Description
Feature Set	Enables or disables the FCoE feature set on Cisco MDS 9000 Family or Nexus 7000 Series switches.
Control	Enables or disables FCoE on Cisco Nexus 5000 Series switches.
Config	Displays the FCoE configuration information on the switch. For example, FC Map and FCF Priority.
VLAN-VSAN Mapping	Displays the VSAN and VLAN IDs with their operational status.

- Step 2** Double-click the relevant FCoE parameter for a switch, and modify the value of the parameter.
- Step 3** In the Information pane toolbar, click the **Apply Changes** icon to save the changes.

Configuring FCoE Using Device Manager

To configure FCoE on a switch using Device Manager, follow these steps:

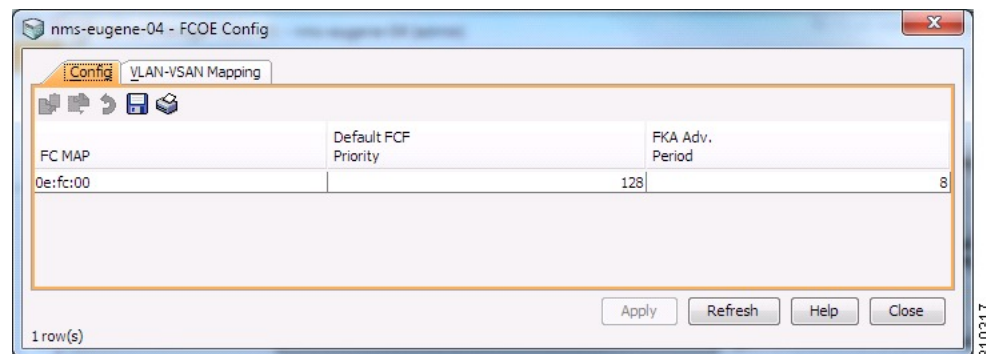
Procedure

- Step 1** Launch Device Manager from the Cisco Nexus 5000 Series switch.
- Step 2** Choose **FCoE > Config**.

You see the FCoE Config dialog box shown in [Figure 53: FCoE Config Dialog Box, on page 447](#).

The Config tab displays the FCoE parameters, such as FC Map, default FCF priority value, and FKA advertisement period, for each Cisco Nexus 5000 Series switch that runs Cisco NX-OS Release 4.1(3) or later releases. [Table 53: FCoE Parameters, on page 447](#) lists the FCoE parameters for a switch.

Figure 53: FCoE Config Dialog Box



- Step 3** Double-click the relevant FCoE parameter for a switch, and modify the value of the parameter.

Step 4 Click **Apply** to save the changes.

Field Descriptions for FCoE

Feature Set

Field	Description
Status	Display only. Displays the FCoE status on the switch.
Command	Display only. Displays the feature set command.
Last Command	Display only. Displays the last feature set command executed on the switch.
Result	Display only. Displays the result of the last feature set command executed on the switch.

Control

Field	Description
Status	Display only. Displays the FCoE status on the switch.
Command	Display only. Displays the feature set command.
Last Command	Display only. Displays the last feature set command executed on the switch.
Result	Display only. Displays the result of the last feature set command executed on the switch.

Config

Field	Description
FC Map	The FCoE Mac Address Prefix used to associate the FCoE Node (ENode).
Default FCF Priority	The default FCoE Initialization Protocol (FIP) priority value advertised by the Fibre Channel Forwarder (FCF) to ENodes.
FKA Adv. Period	The time interval at which FIP Keepalive (FKA) messages are transmitted to the MAC address of the ENode.

VSAN-VLAN Mapping

Field	Description
VSAN Id	The ID of the VSAN.
VLAN Id	The ID of the VLAN.

Field	Description
Oper State	Shows the operational state of this VLAN-VSAN association entry.

Additional References

For additional information related to implementing FCoE, see the following section:

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	Cisco MDS 9000 Family Command Reference

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIB	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

Feature History for FCoE

[Table 54: Feature History for FCoE](#), on page 450 lists the release history for this feature. Only features that were introduced or modified in 5.0(1a) or a later release appear in the table.

Table 54: Feature History for FCoE

Feature Name	Releases	Feature Information
Configuring FCoE	5.2(1)	<p>Added information about discovering Cisco Nexus 7000 and Cisco MDS 9000 Family switches using the FCoE wizard. FICON tape acceleration over FCIP efficiently utilizes the tape device by decreasing idle time.</p> <p>The following sections provide information about this feature:</p>



CHAPTER 18

Configuring Dense Wavelength Division Multiplexing

- [Configuring Dense Wavelength Division Multiplexing, on page 451](#)

Configuring Dense Wavelength Division Multiplexing

Information About DWDM

Dense Wavelength-Division Multiplexing (DWDM) multiplexes multiple optical carrier signals on a single optical fiber. DWDM uses different wavelengths to carry various signals.

To establish a DWDM link, both ends of an Inter Switch Link (ISL) need to be connected with DWDM SFPs (small form-factor pluggable) at each end of the link. To identify a DWDM link, DCNM-SAN discovers the connector type on the Fiber Channel (FC) ports. If the ISL link is associated with the FC ports at each end, then the FC port uses DWDM SFP to connect the links.

Cisco DCNM for SAN discovers FC ports with DWDM SFPs and the ISLs associated with the FC ports. The DCNM-SAN Client displays ISL with DWDM attribute on the topology map.



Note The Fabric Shortest Path First (FSPF) database only displays an ISL link, which is connected with DWDM SFPs at both ends.

Configuring X2 DWDM Transceiver Frequency using DCNM Manager



Note This feature is supported only in MDS 9134 modules. With MDS 9134 modules, the 10-Gigabit Ethernet ports must be in a down state when you configure the X2 transceiver frequency.

To configure the X2 DWDM transceiver frequency using Device Manager, follow these steps:

Procedure

-
- Step 1** From the Device Manager menu bar, select **Physical > Modules**.
The module configuration window is displayed.
- Step 2** Choose an XcvrFrequencyConfig option button.
- Step 3** Click **Apply**.
-

Configuring X2 DWDM Transceiver Frequency using DCNM-SAN

To configure the X2 DWDM transceiver frequency, follow these steps:

Procedure

-
- Step 1** From the Physical Attributes pane, select **Hardware**.
The module configuration window is displayed.
- Step 2** Click the **Card Module Config** tab.
- Step 3** In the X2 XcvrFrequencyConfig column, choose an option.
- Step 4** Click **Apply**.
-

Monitoring DWDM Links

The DCNM-SAN Client displays DWDM links with a “dash-dash” pattern. The tooltip for the link displays “DWDM” to indicate its link type.

To view the DWDM link, follow these steps:

Procedure

-
- Step 1** Select the switch in the Logical Domain region.
- Step 2** Select ISL in the Physical Attributes region.
The Information pane displays the ISL’s information.
- Step 3** Click the **Physical** tab.
You see the ISL in the Information pane.
The ISL’s Physical table displays the connector type as **sfpDwdm**.
Move the mouse over the link to see the tooltip as DWDM indicating the link type.

Step 4 Perform a Dump Discovery of ISL to list all ISLs. DWDM links are listed with [DWDM].

Field Descriptions for DPVM

This section displays the following field descriptions for this feature.

DPVM Actions

Field	Description
Action	Helps in activating the set of bindings.
Result	Indicates the outcome of the activation.
Status	Indicates the state of activation. If true, then activation has been attempted as the most recent operation. If false, then an activation has not been attempted as the most recent operation.
CopyActive to Config	When set to copy(1), results in the active (enforced) binding database to be copied on to the configuration binding database. The learned entries are also copied.
Auto Learn Enable	Helps to learn the configuration of devices logged into the local device on all its ports and the VSANs to which they are associated.
Auto Learn Clear	Assists in clearing the auto-learned entries.
Clear WWN	Represents the port WWN (pWWN) to be used for clearing its corresponding auto-learned entry.

DPVM Config Database

Field	Description
Type	Specifies the type of the corresponding instance of cdpvmLoginDev object.
WWN or Name	Represents the logging in device.
VSAN Id	Represents the VSAN to be associated to the port on the local device on which the device represented by cdpvmLoginDev logs in.
Switch Interface	Represents the device alias.

DPVM Active Database

Field	Description
Type	Specifies the type of the corresponding instance of cdpvmEnfLoginDev.

Field	Description
WWN or Name	Represents the logging in device address.
VSAN Id	Represents the VSAN of the port on the local device through which the device represented by cdpvmEnfLoginDev logs in.
Interface	Represents the device alias.
IsLearnt	Indicates whether this is a learned entry or not. If true, then it is a learned entry. If false, then it is not.

Additional References

For additional information related to implementing VSANs, see the following section:

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	Cisco MDS 9000 Family Command Reference

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To locate and download MIBs, go to the following URL: http://www.cisco.com/dc-os/mibs



CHAPTER 19

Configuring and Managing VSANs

- [Configuring and Managing VSANs, on page 455](#)

Configuring and Managing VSANs

You can achieve higher security and greater stability in Fibre Channel fabrics by using virtual SANs (VSANs) on Cisco MDS 9000 Family switches and Cisco Nexus 5000 Series switches. 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

A VSAN is a virtual storage area network (SAN). A SAN is a dedicated network that interconnects hosts and storage devices primarily to exchange SCSI traffic. In SANs, you use the physical links to make these interconnections. A set of protocols run over the SAN to handle routing, naming, and zoning. You can design multiple SANs with different topologies.

With the introduction of VSANs, the network administrator can build a single topology containing switches, links, and one or more VSANs. Each VSAN in this topology has the same behavior and property of a SAN. A VSAN has the following additional features:

- Multiple VSANs can share the same physical topology.
- The same Fibre Channel IDs (FC IDs) can be assigned to a host in another VSAN, thus increasing VSAN scalability.
- Every instance of a VSAN runs all required protocols such as FSPF, domain manager, and zoning.
- Fabric-related configurations in one VSAN do not affect the associated traffic in another VSAN.
- Events causing traffic disruptions in one VSAN are contained within that VSAN and are not propagated to other VSANs.


This section describes VSANs and includes the following topics:

VSANs Topologies

The switch icons shown in both [Figure 54: Logical VSAN Segmentation](#), on page 456 and [Figure 55: Example of Two VSANs](#), on page 457 indicate that these features apply to any switch in the Cisco MDS 9000 Family.


[Figure 54: Logical VSAN Segmentation](#), on page 456 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 54: Logical VSAN Segmentation



[Figure 55: Example of Two VSANs](#), on page 457 shows a physical Fibre Channel switching infrastructure with two defined VSANs: VSAN 2 (dashed) and VSAN 7 (solid). VSAN 2 includes hosts H1 and H2, application servers AS2 and AS3, and storage arrays SA1 and SA4. VSAN 7 connects H3, AS1, SA2, and SA3.

Figure 55: Example of Two VSANs



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

Without VSANs, a network administrator would need separate switches and links for separate SANs. By enabling VSANs, the same switches and links may be shared by multiple VSANs. VSANs allow SANs to be built on port granularity instead of switch granularity. [Figure 55: Example of Two VSANs , on page 457](#) illustrates that a VSAN is a group of hosts or storage devices that communicate with each other using a virtual topology defined on the physical SAN.

The criteria for creating such groups differ based on the VSAN topology:

- VSANs can separate traffic based on the following requirements:
 - Different customers in storage provider data centers
 - Production or test in an enterprise network
 - Low and high security requirements
 - Backup traffic on separate VSANs
 - Replicating data from user traffic
- VSANs can meet the needs of a particular department or application.

VSAN Advantages

VSANs offer the following advantages:

- **Traffic isolation**—Traffic is contained within VSAN boundaries and devices reside only in one VSAN ensuring absolute separation between user groups, if desired.
- **Scalability**—VSANs are overlaid on top of a single physical fabric. The ability to create several logical VSAN layers increases the scalability of the SAN.
- **Per VSAN fabric services**—Replication of fabric services on a per VSAN basis provides increased scalability and availability.
- **Redundancy**—Several VSANs created on the same physical SAN ensure redundancy. If one VSAN fails, redundant protection (to another VSAN in the same physical SAN) is configured using a backup path between the host and the device.
- **Ease of configuration**—Users can be added, moved, or changed between VSANs without changing the physical structure of a SAN. Moving a device from one VSAN to another only requires configuration at the port level, not at a physical level.

Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

VSANs Versus Zones


You can define multiple zones in a VSAN. Because two VSANs are equivalent to two unconnected SANs, zone A on VSAN 1 is different and separate from zone A in VSAN 2. [Table 55: VSAN and Zone Comparison](#), on page 458 lists the differences between VSANs and zones.

Table 55: VSAN and Zone Comparison

VSAN Characteristic	Zone Characteristic
VSANs equal SANs with routing, naming, and zoning protocols.	Routing, naming, and zoning protocols are not available on a per-zone basis.
—	Zones are always contained within a VSAN. Zones never span two VSANs.
VSANs limit unicast, multicast, and broadcast traffic.	Zones limit unicast traffic.
Membership is typically defined using the VSAN ID to Fx ports.	Membership is typically defined by the pWWN.
An HBA or a storage device can belong only to a single VSAN—the VSAN associated with the Fx port.	An HBA or storage device can belong to multiple zones.
VSANs enforce membership at each E port, source port, and destination port.	Zones enforce membership only at the source and destination ports.
VSANs are defined for larger environments (storage service providers).	Zones are defined for a set of initiators and targets not visible outside the zone.
VSANs encompass the entire fabric.	Zones are configured at the fabric edge.

The following figure shows the possible relationships between VSANs and zones. In VSAN 2, three zones are defined: zone A, zone B, and zone C. Zone C overlaps both zone A and zone B as permitted by Fibre

Channel standards. In VSAN 7, two zones are defined: zone A and zone D. No zone crosses the VSAN boundary—they are completely contained within the VSAN. Zone A defined in VSAN 2 is different and separate from zone A defined in VSAN 7.



VSAN Configuration

VSANs have the following attributes:

- **VSAN ID**—The VSAN ID identifies the VSAN as the default VSAN (VSAN 1), user-defined VSANs (VSAN 2 to 4093), and the isolated VSAN (VSAN 4094).
- **State**—The administrative state of a VSAN can be configured to an active (default) or suspended state. Once VSANs are created, they may exist in various conditions or states.
 - The active state of a VSAN indicates that the VSAN is configured and enabled. By enabling a VSAN, you activate the services for that VSAN.
 - The suspended state of a VSAN indicates that the VSAN is configured but not enabled. If a port is configured in this VSAN, it is disabled. Use this state to deactivate a VSAN without losing the VSAN's configuration. All ports in a suspended VSAN are disabled. By suspending a VSAN, you can preconfigure all the VSAN parameters for the whole fabric and activate the VSAN immediately.
- **VSAN name**—This text string identifies the VSAN for management purposes. The name can be from 1 to 32 characters long and it must be unique across all VSANs. By default, the VSAN name is a concatenation of VSAN and a four-digit string representing the VSAN ID. For example, the default name for VSAN 3 is VSAN0003.



Note A VSAN name must be unique.

- 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.

**Note**

OX ID based load balancing of IVR traffic from IVR-enabled switches is not supported on Generation 1 switching modules. OX ID based load balancing of IVR traffic from a non-IVR MDS 9000 Family switch should work. Generation 2 switching modules support OX ID based load balancing of IVR traffic from IVR-enabled switches.

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.

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—By assigning VSANs to ports.

See the [Assigning Static Port VSAN Membership, on page 464](#).

- Dynamically—By assigning VSANs based on the device WWN. This method is referred to as dynamic port VSAN membership (DPVM).

See [Chapter 25, “Creating Dynamic VSANs.”](#)

Trunking ports have an associated list of VSANs that are part of an allowed list (refer to the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide*).

About the Default VSAN

The factory settings for switches in the Cisco MDS 9000 Family 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. All non-trunking ports are transferred to this VSAN when the VSAN to which they belong is deleted. This avoids an implicit transfer of ports to the default VSAN or to another configured VSAN. All ports in the deleted VSAN are isolated (disabled).



Note When you configure a port in VSAN 4094 or move a port to VSAN 4094, that port is immediately isolated.



Caution Do not use an isolated VSAN to configure ports.



Note Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.


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.

- 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 (refer to the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide*).

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.

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.

About Interop Mode

Interoperability enables the products of multiple vendors to come into contact with each other. Fibre Channel standards guide vendors towards common external Fibre Channel interfaces. See the [“Switch Interoperability” section on page 27-5](#).

About FICON VSANs

You can enable FICON in up to eight VSANs. See the [“FICON VSAN Prerequisites” section on page 24-6](#).

Host Provisioning Wizard

The Host Provisioning wizard provides an intuitive way to commission a new host or decommission an existing host without requiring the use of multiple tools and features. The wizard allows you to create a device alias, and configure DPVM, zoning, and flow creation.

Licensing Requirements for VSAN

The following table shows the licensing requirements for this feature:

License	License Description
ENTERPRISE_PKG	The enterprise license is required to enable VSAN. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .

Default Settings

[Table 56: Default VSAN Parameters](#), on page 462 lists the default settings for all configured VSANs.

Table 56: 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 VSANs

This section includes the following topics:

Multi-tenancy with MDS9000 and DCNM for SAN

Cisco DCNM is capable of providing users a partial view of the installed devices. To ensure that you have a clear understanding for multi-tenancy functionality on DCNM. Following is an example of multi-tenancy applied to a Fiber Channel fabric.

Imagine to have two users configured on your fabric. The first one is the SAN administrator, with full privileges. The second user, called vsanUser, has limited privileges enforced by the Role Based Access Control capabilities of MDS9000 devices and DCNM for SAN. In our example, the user vsanUser is only allowed to see and work on VSAN 2 and 444. He is not allowed to act upon all other VSANs in the fabric. In other words, the user vsanUser has read-write capabilities on VSAN 2 and 444 but he has not even read-only access to other VSANs. This user was configured with the custom role vsanRole as indicated below:

```
sw172-22-46-182# sh role name vsanRole
```

```
Role: vsanRole
```

```
vsan policy: deny
```

```
Permitted vsans: 2,444
```

```
-----  
Rule Type Command-type Feature  
-----
```

```
1. permit show *
```

```
2. permit config *
```

```
3. permit exec *
```

```
sw172-22-46-182#
```

```
vsanUser md5 des(no) vsanRole
```

When user vsanUser belonging to role vsanRole opens a fabric with multiple VSAN via DCNM SAN client, he will only see vsans 2 and 444. Instead, the SAN administrator would see all VSANs configured on switches.



Note

DCNM SAN is not doing any of this filtering; MDS9000 switch is the filtering point enforced by roles.

Creating VSANs


You cannot configure any application-specific parameters for a VSAN before creating the VSAN.

To create and configure VSANs, follow these steps:

Procedure

Step 1

Click the Create VSAN icon.



Note As of Cisco SAN-OS Release 3.1(2) and later, if you check the Static Domain IDs check box, DCNM-SAN 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) Select the **FICON** and **Enable Fabric Binding for Selected Switches** options if you want these features enabled.
- See the *Configuring FICON* section and refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide* for details.
- Step 8** Complete the fields in this dialog box and click Create to add the VSAN or click **Close**.

Assigning Static Port VSAN Membership

To statically assign VSAN membership for an interface, follow these steps:

Procedure

- Step 1** Choose FC Interfaces > Physical from the Physical Attributes pane. You see the interface configuration in the Information pane.
- Step 2** Click the General tab.
You see the Fibre Channel general physical information. Double-click and complete the PortVSAN field.
- Step 3** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

Deleting Static VSANs

To delete a VSAN and its attributes, follow these steps:

Procedure

- Step 1** Select **All VSANs** from the Logical Domains pane.
The VSANs in the fabric are listed in the Information pane.
- Step 2** Right-click the VSAN that you want to delete and select **Delete Row** from the drop-down menu.

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.

Configuring Load Balancing

To configure load balancing on an existing VSAN, follow these steps:

Procedure


- Step 1** Choose Fabric xx > All VSANs from the Logical Domains pane.
You see the VSAN configuration in the Information pane.
- Step 2** Select a VSAN and complete the LoadBalancing field.
- Step 3** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

Commissioning a Host

To commission a new host, follow these steps:

Procedure

- Step 1** From the DCNM-SAN window, select **Tools > Host Provisioning**.
The Host Provisioning wizard window is displayed.
- Step 2** Click the **Commission** radio button.
- Step 3** Click [...] and select the host from the existing configurations or VSAN, or enter the WWN of a host that is not in VSAN or not configured yet.



If the host configuration already exists, the switch, device alias, and VSAN information are populated in the window.

If the configuration does not exist already, enter a device alias for the WWN, enter a switch where the configuration will be initiated, and select a VSAN to which the host should belong. The entries are created and saved when you click Next in the Host Provisioning wizard window.

- Step 4** Uncheck the **Skip Zoning** check box.
- Step 5** Click Next. The Select Targets and the Select Zone windows appear.

Step 6 Uncheck the **Skip DPVM** check box.

Step 7 Click **Next**. The DPVM entries are created.

Step 8 Click **Next**.

The Select Targets window appears.

Note The Host Provisioning wizard requires that basic and enhanced device alias, DPVM, and CFS to be enabled in all switches in the selected VSAN.

Step 9 Select the target with which the host needs to communicate, and click **Add**.

The target entry is moved to the bottom of the window.

Step 10 Click **Next**.

The Select Zone window appears.

Step 11 Select a zone and check the **Create Flow after Activation** check box.

The host and storage are added to a zone and the zone is activated, and a flow between host and storage is created when you click **Finish**.

Step 12 Click **Finish**.

The device alias and DPVM entries are created, a zone is created and activated, and the flow is created based on the check boxes you checked.

Decommissioning a Host

To decommission an existing host, follow these steps:

Procedure

Step 1 From the DCNM-SAN window, select **Tools > Host Provisioning**.

The Select Host window appears.

Step 2 Click the **Decommission** radio button.

Step 3 Click [...] and select the host from the existing configurations or VSAN, or enter the WWN of a host that is not in VSAN.

The device alias and DPVM state from all of the switches in the selected VSAN are populated if device alias with CFS and CFS DPVM are enabled and if the WWN is an eight-byte number.

Step 4 Click **Finish**. The device aliases are removed.

Step 5 Uncheck the **Skip Zoning** check box.

The WWN zone member is removed from all zones. If the zones without a WWN member become single member zones, these zones also are removed.

Step 6 Click **Finish**. If there is a local active zone set change due to the removal of zones, the appropriate zone set is activated.

Step 7 Uncheck the **Skip DPVM** check box.

- Step 8** Click **Finish**. The DPVM entry is removed.
- Step 9** Click **Next**.
The Decommission Zones window appears.
- Step 10** Check the **Remove Flow after Deactivation** check box.
The flow entry associated with the host is removed when you click **Finish**.
- Step 11** Click **Finish**.
The device alias and DPVM entries are deleted, the zone is deactivated and deleted (if it has only one member after removing the host), and the flow is deleted depending on the check boxes you checked.

Displaying Isolated VSAN Membership

The **show vsan 4094 membership** command displays all ports associated with the isolated VSAN.

To display interfaces that exist in the isolated VSAN, follow these steps:

Procedure

- Step 1** Expand Fabric xx, and then select All VSANs in the Logical Domains pane.
You see the VSAN configuration in the Information pane.
- Step 2** Click the Isolated Interfaces tab.
You see the interfaces that are in the isolated VSAN.

Field Descriptions for VSAN

The following are the field descriptions for VSAN.

VSAN General

Field	Description
Name	The name of the VSAN. Note that default value will be the string VSANxxxx where xxxx is value of vsanIndex expressed as 4 digits. For example, if vsanIndex is 23, the default value is VSAN0023.
Mtu	The MTU of the VSAN. Normally, this is 2112.
LoadBalancing	The type of load balancing used on this VSAN. <ul style="list-style-type: none"> • srcdst— use source and destination ID for path selection • srcdst 0xld— use source, destination, and exchange IDs

Field	Description
InterOp	The interoperability mode of the local switch on this VSAN. <ul style="list-style-type: none"> • standard • interop-1 • interop-2 • interop-3
AdminState	The state of this VSAN.
OperState	The operational state of the VSAN.
InOrderDelivery	The InorderDelivery guarantee flag of device. If true, then the inorder delivery is guaranteed. If false, it is not guaranteed.
DomainId	Specifies an insistent domain ID.
FICON	True if the VSAN is FICON-enabled.
Network Latency	Network latency of this switch on this VSAN. This is the time interval after which the frames are dropped if they are not delivered in the order they were transmitted.

VSAN Membership

Field	Description
Switch	Name of the switch
Ports	FC ports in VSAN
Channels	PortChannels in VSAN
FCIP	FCIP Interfaces in VSAN
iSCSI	iSCSI Interfaces in VSAN
FICON	Interfaces in VSAN by FICON
FC Virtual Interface	Virtual FC interfaces in VSAN

VSAN Interop-4 WWN

Field	Description
VSAN ID	The ID of the VSAN containing the McData switch.
WWN	The WWN of the McData switch.

VSAN Timers

Field	Description
VSAN Id	The ID of the VSAN.
R_A_TOV	The Resource_Allocation_Timeout Value used for FxPorts as the timeout value for determining when to reuse an NxPort resource such as a Recovery_Qualifier. It represents E_D_TOV plus twice the maximum time that a frame may be delayed within the fabric and still be delivered. Note that all switches in a fabric should be configured with the same value of this timeout.
D_S_TOV	The Distributed_Services_Timeout Value which indicates that how long a distributed services requestor will wait for a response.
E_D_TOV	The Error_Detect_Timeout Value used for FxPorts as the timeout value for detecting an error condition. Note that all switches in a fabric should be configured with the same value of this timeout. Note that value must be less than value of D_S_TOV.
NetworkDropLatency	Network latency of this switch on this VSAN.

VSAN Default Zone Policies

Field	Description
Zone Behavior	Represents the initial value for default zone behavior on a VSAN when it is created. If a VSAN were to be deleted and re-created again, the default zone behavior will be set to the value specified for this object.
Propagation Mode	Represents the initial value for zone set propagation mode on a VSAN when it is created. If a VSAN were to be deleted and re-created again, the zone set propagation mode will be set to the value specified for this object.



CHAPTER 20

Discovering SCSI Targets

- [Discovering SCSI Targets, on page 471](#)

Discovering SCSI Targets

This chapter describes the SCSI LUN discovery feature provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

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 an NMS can access this information.
- To report device capacity, serial number, and device ID information.
- To register the initiator and target features with the name server.

The SCSI LUN discovery feature uses the local domain controller Fibre Channel address. It uses the local domain controller as the source FC ID, and performs SCSI INQUIRY, REPORT LUNS, and READ CAPACITY commands on SCSI devices.

The SCSI LUN discovery feature is initiated on demand, through CLI or SNMP. This information is also synchronized with neighboring switches, if those switches belong to the Cisco MDS 9000 Family.

This section includes the following topics:

About Starting SCSI LUN Discovery

SCSI LUN discovery is done on demand.

Only Nx ports that are present in the name server database and that are registered as FC4 Type = SCSI_FCP are discovered.

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.

Use the **custom-list** option to initiate this discovery.

Licensing Requirements for SCSI

The following table shows the licensing requirements for this feature:

Feature	License Requirement
ENTERPRISE_PKG	The enterprise license is required to enable the SCSI flow statistics. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .
FM_SERVER_PKG	The Cisco DCNM for SAN Package is required to enable the traffic analyzer for SCSI flow statistics. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .

Discovering SCSI Targets

This section includes the following topics:

Starting SCSI LUN Discovery using Device Manager

To begin SCSI LUN discovery using Device Manager, follow these steps:

Procedure

-
- Step 1** Choose **FC > Advanced > LUNs**.
You see the 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.
-

Initiating Customized Discovery using Device Manager

To initiate a customized discovery using Device Manager, follow these steps:

Procedure

-
- Step 1** From the **VSAN** drop-down menu, select the VSAN in which you want to initiate a customized discovery.
 - Step 2** Click **FC > Advanced > LUNs**.
You see the LUN Configuration dialog box.
 - Step 3** Set **StartDiscovery** to **local**, **remote** or **both**.
 - Step 4** Fill in the **DiscoveryType** and **OS** fields.

Step 5 Click **Apply** to begin discovery.

Field Descriptions for SCSI Targets

The following are the field descriptions for SCSI targets.

iSCSI Connection

Field	Description
LocalAddr	The local Internet network address used by this connection.
RemoteAddr	The remote Internet network address used by this connection.
CID	The iSCSI connection ID for this connection.
State	The current state of this connection, from an iSCSI negotiation point of view. <ul style="list-style-type: none"> • login— The transport protocol connection has been established, but a valid iSCSI login response with the final bit set has not been sent or received. • full—A valid iSCSI login response with the final bit set has been sent or received. • logout— A valid iSCSI logout command has been sent or received, but the transport protocol connection has not yet been closed.
MaxRecvDSLen	The maximum data payload size supported for command or data PDUs in use within this connection. The size is reported in bytes even though the negotiation is in 512K blocks.
SendMarker	Indicates whether or not this connection is inserting markers in its outgoing data stream.
HeaderDigest	The iSCSI header digest scheme in use within this connection.
DataDigest	The iSCSI data digest scheme in use within this connection.

iSCSI Initiators

Field	Description
Name or IP Address	A character string that is a globally unique identifier for the node represented by this entry.
VSAN Membership	The list of configured VSANs the node represented by this entry can access.
Dynamic	If true, then the node represented by this entry is automatically discovered.
Initiator Type	Indicates whether the node is a host that participates in iSCSI load-balancing.
Persistent Node WWN	If true, then the same FC address is assigned to the node if it were to be represented again in the FC domain with the same node name. Note that the node FC address is either automatically assigned or manually configured.
SystemAssigned Node WWNN	If true, the FC address is automatically assigned to this node. If false, then the FC address has to be configured manually.

Field	Description
Node WWN	The persistent FC address of the node.
Persistent Port WWN	If true, then the same FC address is assigned to the ports of the node if it were to be represented again in the FC domain with the same node name.
Port WWN	All the FC port addresses associated with this node.
AuthUser	This is the only CHAP user name that the initiator is allowed to log in with.
Target UserName	(Optional) The user name to be used for login. If you do not supply a username, the global user name is used.
Target Password	(Optional) The password to be used for login. If you do not supply a password, the global password is used.
Load Metric	A configured load metric of this iSCSI initiator for the purpose of iSCSI load balancing.
Auto Zone Name	The zone name that is used when the system creates automatic zone for this initiator's specific list of targets.

iSCSI Targets

Field	Description
Dynamically Import FC Targets	Check this option to dynamically import FC targets into the iSCSI domain. A target is not imported if it already exists in the iSCSI domain.
iSCSI Name	The iSCSI name of the node represented by this entry.
Dynamic	Indicates if the node represented by this entry was either automatically discovered or configured manually.
Primary Port WWN	The FC address for this target.
Secondary Port WWN	The optional secondary FC address for this target. This is the FC address used if the primary cannot be reached.
LUN Map iSCSI	The configured default logical unit number of this LU.
LUN Map FC Primary	The logical unit number of the remote LU for the primary port address.
LUN Map FC Secondary	The logical unit number of the remote LU for the secondary port address.
Initiator Access All	If true, then all the initiators can access this target even those which are not in the initiator permit list of this target. If false, then only initiators which are in the permit list are allowed access to this target.
Initiator Access List	Lists all the iSCSI nodes that are permitted to access the node represented by this entry. If AllAllowed is false and the value of List is empty, then no initiators are allowed to access this target.

Field	Description
Advertised Interfaces	Lists all the interfaces on which the target could be advertised.
Trespass Mode	The trespass mode for this node. Every iSCSI target represents one or more port(s) on the FC target. If true, the node instructs the FC node to present all LUN I/O requests to secondary port if the primary port is down.
RevertToPrimaryPort	Indicates if it is required to revert back to primary port if the FC target comes back online.

iSCSI Session Initiators

Field	Description
Name or IP Address	The name or IP address of the initiator port.
Alias	The initiator alias acquired at login.

iSCSI Global

Field	Description
AuthMethod	The authentication method.
InitiatorIdleTimeout	The time for which the gateway (representing a FC target) waits from the time of last iSCSI session to a iSCSI initiator went down, before purging the information about that iSCSI initiator.
iSLB ZonesetActivate	Checking this option performs automatic zoning associated with the initiator targets
DynamicInitiator	This field determines how dynamic iSCSI initiators are created. Selecting the iSCSI option (default) creates dynamic iSCSI initiators. If you select iSLB then the an iSLB dynamic initiator is created. Selecting the deny option does not allow dynamic creation of the initiators.
Target UserName	The default user name used for login. If an initiator user name is specified, that user name is used instead.
Target Password	The default password used for login. If an initiator password is specified, that password is used instead.

iSCSI Session Statistics

Field	Description
PDU Command	The count of Command PDUs transferred on this session.
PDU Response	The count of Response PDUs transferred on this session.
Data Tx	The count of data bytes that were transmitted by the local iSCSI node on this session.

Field	Description
Data Rx	The count of data bytes that were received by the local iSCSI node on this session.
Errors Digest	Authentication errors.
Errors CxnTimeout	Connection timeouts.

iSCSI iSLB VRRP

Field	Description
VrId, IpVersion	The virtual router number and the IP version (IPv4, IPv6, or DNS).
Load Balance	Indicates whether load balancing is enabled.

iSCSI Initiator Access

Field	Description
Initiator Name	The iSCSI node name.

iSCSI Initiator PWWN

Field	Description
Port WWN	The FC address for this entry.

iSCSI Sessions

Field	Description
Type	Type of iSCSI session: <ul style="list-style-type: none"> • normal—Session is a normal iSCSI session • discovery—Session is being used only for discovery.
TargetName	If Direction is Outbound, this will contain the name of the remote target.
Vsan ID	The VSAN to which this session belongs to.
ISID	The initiator-defined portion of the iSCSI session ID.
TSIH	The target-defined identification handle for this session.

iSCSI Sessions Detail

Field	Description
ConnectionNumber	The number of transport protocol connections that currently belong to this session.
ImmediateData	Whether the initiator and target have agreed to support immediate data on this session.
Initial	If true, the initiator must wait for a Ready-To-Transfer before sending to the target. If false, the initiator may send data immediately, within limits set by FirstBurstSize and the expected data transfer length of the request.
MaxOutstanding	The maximum number of outstanding Ready-To-Transfers per task within this session.
First	The maximum length supported for unsolicited data sent within this session.
Max	The maximum number of bytes which can be sent within a single sequence of Data-In or Data-Out PDUs.
Sequence	If false, indicates that iSCSI data PDU sequences may be transferred in any order. If true indicates that data PDU sequences must be transferred using continuously increasing offsets, except during error recovery.
PDU	If false, iSCSI data PDUs within sequences may be in any order. If true indicates that data PDUs within sequences must be at continuously increasing addresses, with no gaps or overlay between PDUs.



CHAPTER 21

Configuring SAN Device Virtualization

- [Configuring SAN Device Virtualization, on page 479](#)

Configuring SAN Device Virtualization

This chapter describes how to configure virtual devices to represent physical end devices for switches running Cisco MDS SAN-OS Release 3.1(2) and later, or NX-OS Release 4.1(1a) and later.

Cisco SAN device virtualization (SDV) is a licensed feature included in the Cisco MDS 9000 Family Enterprise package (ENTERPRISE_PKG). Refer to the *Cisco NX-OS Family Licensing Guide* for details about acquiring licenses.

Information About SDV

As of Cisco SAN-OS Release 3.1(2) and later, you can use Cisco SAN device virtualization to create virtual devices that represent physical end-devices. Virtualization of SAN devices accelerates swapout or failover to a replacement disk array, and it also minimizes downtime when replacing host bus adapters (HBAs) or when rehosting an application on a different server.

SAN device virtualization enables you to:

- Reduce the amount of time it takes for data migration, and ultimately the overall amount of downtime.
- Improve ease-of-use and reduce the possibility of user-introduced errors during the failover by performing the operation in a single step.
- Easily scale to larger numbers of targets.

SAN devices that are virtualized can be either initiators or targets. You can virtualize targets to create a *virtual target* and also virtualize initiators to create a *virtual initiator*. SAN device configurations do not distinguish between virtual initiators and virtual targets (see [Figure 56: Target Virtualization, on page 480](#) and [Figure 57: Initiator Virtualization, on page 480](#)).

Figure 56: Target Virtualization






Figure 57: Initiator Virtualization

**Note**

While most of the examples in this chapter describe target virtualization, the initiator virtualization functions similarly.

Typically, today's deployments for handling device failures are designed for high availability (HA), with redundancy being a key part of this design. Consider the situation where a target is designed to be redundant. Two arrays are deployed—a primary and secondary in this situation. Enterprises often use some type of consistency technology (such as EMF SRDF) between the primary and secondary arrays to ensure that the secondary is a mirrored copy of the production LUN. However, if the primary array fails, it must be replaced by the secondary because all I/O must occur on the secondary array. Problems can occur because the time required to bring the secondary array up and have it working often takes longer than most can afford ([Figure 58: Typical Deployment for Handling Device Failures Before SDV](#), on page 480 illustrates this dilemma).

Figure 58: Typical Deployment for Handling Device Failures Before SDV



If a storage array is replaced without using Cisco SDV, then it may require the following actions:

- Taking down a server to modify zoning and account for the new array.
- Changing the Cisco NX-OS configuration to accommodate Fibre Channel IDs (FC IDs) and pWWNs of the new array.
- Changing a server configuration to accommodate the new FC IDs and pWWNs.

More specifically, without SDV you might experience the following conditions:

- It can take a considerable amount of time to configure a secondary device for a typical production environment.
- In the zoning configuration, all the initiators must be rezoned with the secondary device, and certain initiators must also be reconfigured. For example, the WWN and FC ID of the secondary device are different, so driver files must be changed and the server must be rebooted.
- Clustering (multiple initiators) compounds the problem, and the failover procedure must be repeated for each server of the cluster. Think of a server cluster as a set of HBAs—any storage array FC ID changes must be performed for each HBA.

SDV enables you to achieve the following performance targets:


- Reduce the amount of time it takes for data migration, and ultimately the overall amount of downtime.
- Easily scale to larger numbers of devices.

[Figure 59: SDV Example, on page 482](#) illustrates the benefits of SDV. In this configuration, disk array Y replaces disk array X. When disk array X was deployed, the user created virtual devices for all the Fibre Channel interfaces using SDV. After data replication from disk array X was completed, the user briefly pauses activity on the application server and relinked disk array Y to the virtual devices used by the server, completing the swapout of disk array X. No zoning changes or host operating system configuration changes were required during the time-critical period when the swap was performed; this significantly minimized application downtime.

**Note**

The array administrator will likely have to perform actions on array Y for it to become a primary device and accept server logins before linking the virtual device to the array Y pWWN.

Figure 59: SDV Example



Key Concepts

The following terms are used throughout this chapter:

- **Virtual device**—The virtualized or proxy representation of the real device, which is registered with the name server and has a pWWN and FC ID. A virtual device exists as long as its real (physical) counterpart is online. The virtual device pWWN and FC ID must be unique and cannot clash with any real device pWWNs and FC IDs.
- **Virtual domain**—Reserved by SDV to assign FC IDs to virtual devices. If the switch that reserved the domain goes down, another switch takes over its role using the same domain.
- **Primary device**—The device that is configured as primary. By default, the primary device becomes the active device if it is online.
- **Secondary device**—The additional device that is configured. By default, the secondary device is standby.
- **Active device**—The device that is currently virtualized is called the active device. By default, the primary device becomes the active device if it is online. The active device is indicated by a (*) symbol.

Automatic Failover and Fallback

As of Cisco MDS NX-OS Release 4.1(1a), SAN device virtualization supports automatic failover and fallback configurations for the virtual devices. In all of the earlier releases, when there was a failure, you needed to manually configure the device as primary to make it active. With the introduction of automatic failover and fallback configurations, the active device is distinguished from the primary device indicated by a (*) symbol.

- **Auto failover**—When there is a failure, the failover auto attribute automatically shuts down the primary device and brings up the secondary device to active state. When the primary device comes back online, it requires user intervention to switchover.
- **Auto failover with fallback**—In addition to automatic failover, when the primary device comes back online after a failover, the primary device is brought to active state and the secondary device moved to standby state.

To configure automatic failover, use the **attribute failover auto** command in SAN device virtualization configuration mode. To configure automatic failover and fallback, use the **attribute failover auto fallback** command. To identify the active device, use the **show sdv database** command.

Resolving Fabric Merge Conflicts

Whenever two fabrics merge, SDV merges its database. A merge conflict can occur when there is a run-time information conflict or configuration mismatch. Run-time conflicts can occur due to:

- Identical pWWNs have been assigned to different virtual devices.
- The same virtual devices are assigned different pWWNs.
- The virtual device and virtual FC ID are mismatched.

A *blank commit* is a commit operation that does not contain configuration changes, and enforces the SDV configuration of the committing switch fabric-wide. A blank commit operation resolves merge conflicts by pushing the configuration from the committing switch throughout the fabric, which reinitializes the conflicting virtual devices. Exercise caution while performing this operation, as it can easily take some virtual devices offline.

Merge failures resulting from a pWWN conflict can cause a failure with the device alias as well. A blank commit operation on a merge-failed VSAN within SDV should resolve the merge failure in the device alias.

You can avoid merge conflicts due to configuration mismatch by ensuring that:

- The pWWN and device alias entries for a virtual device are identical (in terms of primary and secondary).
- There are no virtual device name conflicts across VSANs in fabrics.

Licensing Requirements for SAN Device Virtualization

The following table shows the licensing requirements for this feature:

License	License Description
ENTERPRISE_PKG	The enterprise license is required for SAN device virtualization. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .

Guidelines and Limitations

As of MDS NX-OS Release 4.1(1a), the following conditions must be considered when configuring the virtual device failover attributes:

- The attribute configuration is supported only with MDS NX-OS Release 4.1(1a) and later. In a mixed mode fabric where earlier releases are combined, the attribute configuration will fail.
- When the failover attribute is configured, if the primary device is offline then the secondary device becomes active.
- When the failover attribute is deleted after the primary device failover to the secondary device, then the primary becomes active if the primary device is online. If the primary device is not online, then the SDV virtual device is shut down.

**Note**

The SDV attributes configuration is supported in Cisco DCNM for SAN Release 4.1(2) and later.

SDV Requirements and Guidelines

Be aware of the following requirements and guidelines as you plan and configure SDV:

- SDV should be enabled on switches where devices that are part of SDV zones are connected.
- SDV does not work for devices connected to non-MDS switches.
- Broadcast zoning is not supported for a zone with a virtual device.
- IVR and SDV cannot be used for the same device. A SDV-virtualized device cannot be part of an IVR zone or zoneset.
- Virtual device names should be unique across VSANs because they are registered with the device alias server, which is unaware of VSANs. For example, if you have enabled SDV and have registered a name, vt1 in both VSAN 1 and VSAN 2, then the device alias server cannot store both entries because they have the same name.
- You cannot specify the same primary device for different virtual devices.
- SDV does not work with soft zoning (*Soft zoning* means that 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; it does not work with the **zone default-zone permit vsan** operation (which would otherwise permit or deny traffic to members in the default zone).
- If devices are not already zoned with the initiators, then you can configure SDV virtual device zones with no negative impact. If they are already zoned, then zoning changes are required.
- The real device-virtual device zone cannot coexist with the real device-real device zone. If the real devices are not already zoned together, then you can configure the real device-virtual device zone with no negative impact. If these devices are already zoned, then adding the real device-virtual device zone may cause the zone activation to fail. If this occurs, then you must delete one of the zones before activation.

For example, a user attempts to create a configuration with zone A, which consists of I, the initiator, and T, the target (I,T), and zone B, which consists of a virtual initiator, VI, and real target, T (zone VI, T). Such a configuration would fail. Likewise, an attempt to configure zone C, which consists of an initiator, I, and target T, with zone D, which consists of an initiator, I, and virtual target, VT (zone I, VT), would also fail.



Caution

There must be at least one SDV-enabled switch that is not a Cisco MDS 9124 Switch between the server and the device that are being virtualized. SDV does not work when initiators and primary devices are connected to the same Cisco MDS 9124 Switch.

Guidelines for Downgrading SDV

As of MDS NX-OS Release 4.1(1a), SDV supports failover and fallback attribute configuration. Downgrading to an earlier release requires you to remove the attribute configurations before downgrading.

As of SAN-OS Release 3.1(3), SDV supports virtual initiators and LUN zoning. Consequently, in SAN-OS Releases 3.1(3) and later, if virtual initiators are configured or SDV devices are configured as LUN-based members of a zone, a configuration check will indicate that downgrading to SAN-OS Release 3.1(2) may be disruptive and is therefore not recommended.

Downgrading with Attributes Configured

As of MDS NX-OS Release 4.1(1a), SDV supports failover and fallback attribute configuration. To successfully downgrade to an earlier release, you must remove the attribute configurations before downgrading.

Downgrading with Virtual Initiators Configured

If SDV virtual initiators are configured, you will be unable to downgrade to SAN-OS Release 3.1(2).

This incompatibility only warns before a downgrade. We recommend that you remove the virtual initiator configuration or shut down the initiator port so that there are no inconsistencies in the downgraded version.



Note We also recommend that you trigger a manual discovery on all the switches before configuring the virtual initiators; in fact, you can trigger the discovery before proceeding to the downgrade by entering the following commands: switch# **discover scsi-target local os all**discovery startedswitch# **discover scsi-target remote os all**discovery startedswitch#

Downgrading with SDV LUN Zoning Configured

The following are downgrade scenarios when SDV LUN zoning is configured:

- Real initiator and SDV virtual target with LUN
- SDV virtual initiator and real target with LUN
- SDV virtual initiator and SDV virtual target with LUN

In each of these cases, a configuration check is registered to prevent users from downgrading to SAN-OS Release 3.1(2). This incompatibility will be disruptive if you proceed with the downgrade.

To avoid the configuration check, delete all the LUN zone members from SDV zones, and then activate the zone set before the downgrade.

Default Settings

Table 57: Default SDV Configuration Parameters , on page 485 lists the default settings for SDV parameters.

Table 57: Default SDV Configuration Parameters

Parameters	Default
enable	disabled

Configuring SDV

SDV is a distributed service and uses Cisco Fabric Services (CFS) distribution to synchronize the databases. When you configure SDV, it starts a CFS session and locks the fabric. When a fabric is locked, Cisco NX-OS software does not allow any configuration changes from a switch other than the switch holding the lock and issues a message to inform users about the locked status. Configuration changes are held in a pending database for the application. You must perform a commit operation to make the configuration active and to release the lock for all switches. You can discard or stop changes from being distributed by entering the **abort** or **clear** command.

Refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide* for more details about CFS.



Note When you enable SDV, CFS distribution is also enabled; CFS distribution cannot be disabled for SDV.

Configuring a Virtual Device

A virtual device is identified by an alphanumeric name of up to 32 characters and defines all the real devices (one primary and one or more secondary) that it represents. Upon the successful creation of a virtual device, the virtual device name is internally registered as the device alias name with the device alias database; the pWWN is automatically assigned by the system using Cisco Organizational Unique Identifier (OUI). A virtual device appears as a real, physical device. You can enumerate up to 128 devices for a virtual device. There is a limit of 4095 on the number of virtual devices that you can create in a single VSAN.



Note As of Cisco MDS SAN-OS Release 3.1(2) and NX-OS Release 4.1(1a), SDV supports up to 1024 virtual devices per VSAN.


To configure a virtual device and commit it to the fabric configuration, follow these steps:

Procedure

- Step 1** Expand SAN in the Logical Domains pane, and then expand the fabric in which your VSAN resides.
- Step 2** Expand the VSAN in which you want to create the virtual target and select SDV. You see the switches in the VSAN that you selected listed in the Information pane.
- Step 3** In the **Control** tab, select enable from the drop-down menu in the Command column to enable SAN device virtualization for a particular switch in the VSAN.
- Step 4** Click the **Apply Changes** icon to commit the configuration change.
- Step 5** Click the **CFS** tab. Confirm that the SAN device virtualization feature is enabled for the switch.
- Step 6** Click the **Virtual Devices** tab and then click the **Create Row** icon.
You see the Create Virtual Devices dialog box.
- Step 7** Select the Virtual Device ID from the drop-down list (ranges from 1 to 4096).
- Step 8** Enter a Name for the Virtual Device. Select the Virtual Domain and enter a Virtual FC ID for the virtual target.
- Step 9** Check only the autoFailover check box or check the autoFailover and primFallback check boxes. For more information, see the [Automatic Failover and Fallback, on page 482](#). You can also change the option in the Option column of the Virtual Devices tab.
- Step 10** Click **Create** to create the virtual target.
- Step 11** Click the **CFS** icon to commit and distribute the configuration changes.

Example - Configuring a Zone for a Virtual Device

The figure below shows a virtual device-name device alias (vt1) zoned with the real devices activated; the primary device is online.



SDV is enabled on a switch and a virtual device is defined. SDV assigns a pWWN for the virtual device, and it is zoned based on the pWWN in a zone. If you later disable SDV, this configuration is lost. If you reenables SDV and create the virtual device using the same name, there is no guarantee that it will get the same pWWN again. You would have to rezone the pWWN-based zone. However, if you perform zoning based on the device-alias name, there are no configuration changes required if or when the pWWN changes. Be sure you understand how device alias modes work before enabling them. Refer to *Distributing Device Alias Services*, for details and requirements about device alias modes.

Linking a Virtual Device with a Physical Device

To link a virtual target with a physical target, follow these steps:

Before you begin

As of MDS NX-OS Release 4.1(1a), the following condition must be considered before linking a device:

If you link to the secondary device which is currently active because of failover, the primary tag is moved to the secondary device and the secondary device becomes the primary device.

Procedure

- Step 1** Click the Real Devices tab and then click the Create Row icon.
- Step 2** Select the Virtual Device ID from the pull-down list or enter an existing ID for the virtual target that you are linking with a physical target.
- Step 3** Select the Real Device ID of the physical target that you are linking with the virtual target.
- Step 4** Click either the pWWN or deviceAlias radio button, and select the appropriate pWWN or device alias from the pull-down menu. The Name field is automatically populated when you select the pWWN or device alias.
- Step 5** Click either the primary or secondary radio button for the Map Type.

Step 6 Click the **CFS** icon to save and distribute these changes, or click Close to discard any unsaved changes.

Field Descriptions for SDV

This section displays the field descriptions for this feature.

SDV Virtual Devices

Field	Description
Name	Represents the name of this virtual device.
Virtual Domain	The user preference for a persistent Domain ID for this virtual device to indicate a specific partition (domain) of the fabric that this virtual device should belong to.
Virtual FCID	The user preference for a persistent FCID for this virtual device.
Port WWN	The assigned pWWN for this virtual device. The agent assigns this value when the configuration is committed.
Node WWN	The assigned nWWN for this virtual device. The agent assigns this value when the configuration is committed.
Assigned FCID	The assigned FCID of this virtual device. The agent assigns this value when the configuration is committed and the real device that this virtual device virtualizes is online.
Real Device Map List	The set of real device(s) that this virtual device virtualizes in this VSAN.

SDV Real Devices

Field	Description
Type	The type of real device identifier represented by the value of the corresponding instance of cFcSdvVirtRealDeviceId that this virtual device virtualizes to.
Name	Represents a real device(s) identifier that this virtual device virtualizes.
Map Type	The mapping association type of the real device(s) (initiator/target).

Additional References

For additional information related to implementing VSANs, see the following section:

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	Cisco MDS 9000 Family Command Reference

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none"> CISCO-FC-SDV-MIB 	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html



CHAPTER 22

Configuring Fibre Channel Write Acceleration

- [Configuring Fibre Channel Routing Services and Protocols, on page 491](#)

Configuring Fibre Channel Routing Services and Protocols

This chapter describes Fibre Channel routing services and protocols.

This chapter includes the following sections:

Information About FSPF

Fabric Shortest Path First (FSPF) is the standard path selection protocol used by Fibre Channel fabrics. The FSPF feature is enabled by default on all Fibre Channel switches. Except in configurations that require special consideration, you do not need to configure any FSPF services. FSPF automatically calculates the best path between any two switches in a fabric. FSPF provides these features:

- Dynamically compute routes throughout a fabric by establishing the shortest and quickest path between any two switches.
- Select an alternative path in the event of the failure of a given path. FSPF supports multiple paths and automatically computes an alternative path around a failed link. It provides a preferred route when two equal paths are available.

FSPF is the protocol currently standardized by the T11 committee for routing in Fibre Channel networks. The FSPF protocol has the following characteristics and features:

- Supports multipath routing.
- Bases path status on a link state protocol.
- Routes hop by hop, based only on the domain ID.
- Runs only on E ports or TE ports and provides a loop free topology.
- Runs on a per VSAN basis. Connectivity in a given VSAN in a fabric is guaranteed only for the switches configured in that VSAN.
- Uses a topology database to keep track of the state of the links on all switches in the fabric and associates a cost with each link.
- Guarantees a fast reconvergence time in case of a topology change. Uses the standard Dijkstra 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.

This section includes the following topics:

FSPF Global Configuration

By default, FSPF is enabled on switches in the Cisco MDS 9000 Family.

Some FSPF features can be globally configured in each VSAN. By configuring a feature for the entire VSAN, you do not have to specify the VSAN number for every command. This global configuration feature also reduces the chance of typing errors or other minor configuration errors.



Note FSPF is enabled by default. Generally, you do not need to configure these advanced features.



Caution The default for the backbone region is 0 (zero). You do not need to change this setting unless your region is different from the default. If you are operating with other vendors using the backbone region, you can change this default to be compatible with those settings.

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 Record Defaults

Each time a new switch enters the fabric, a link state record (LSR) is sent to the neighboring switches, and then flooded throughout the fabric. [Table 58: LSR Default Settings](#), on page 492 displays the default settings for switch responses.

Table 58: 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.

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.

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.

About Dead Time Intervals

You can set the FSPF dead time interval to specify the maximum interval for which a hello message must be received before the neighbor is considered lost and removed from the database. The integer value can range from 1 to 65,535 seconds.



Note This value must be the same in the ports at both ends of the ISL.



Caution An error is reported at the command prompt if the configured dead time interval is less than the hello time interval.

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.

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.


FSPF Routes

FSPF routes traffic across the fabric, based on entries in the FSPF database. These routes can be learned dynamically, or configured statically.

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 60: Fibre Channel Routes](#), on page 494).

Figure 60: Fibre Channel Routes



Note

Other than in VSANs, runtime checks are not performed on configured and suspended static routes.

About Broadcast and Multicast Routing

Broadcast and multicast in a Fibre Channel fabric uses the concept of a distribution tree to reach all switches in the fabric.

FSPF provides the topology information to compute the distribution tree. Fibre Channel defines 256 multicast groups and one broadcast address for each VSAN. Switches in the Cisco MDS 9000 Family only use broadcast routing. By default, they use the principal switch as the root node to derive a loop-free distribution tree for multicast and broadcast routing in a VSAN.



Caution

All switches in the fabric should run the same multicast and broadcast distribution tree algorithm to ensure the same distribution tree.

To interoperate with other vendor switches (following FC-SW3 guidelines), the Cisco SAN-OS and Cisco NX-OS Release 4.1(1b) and later releases uses the lowest domain switch as the root to compute the multicast tree in interop mode.

About Multicast Root Switch

By default, the **native** (non-interop) mode uses the principal switch as the root. If you change the default, be sure to configure the same mode in all switches in the fabric. Otherwise, multicast traffic could encounter potential loop and frame-drop problems.



Note

The operational mode can be different from the configured interop mode. The interop mode always uses the lowest domain switch as the root.

Use the **mcast root lowest vsan** command to change the multicast root from the principal switch to lowest domain switch.

In-Order Delivery

In-order delivery (IOD) of data frames guarantees frame delivery to a destination in the same order that they were sent by the originator.

Some Fibre Channel protocols or applications cannot handle out-of-order frame delivery. In these cases, switches in the Cisco MDS 9000 Family preserve frame ordering in the frame flow. The source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) identify the flow of the frame.

On any given switch with IOD enabled, all frames received by a specific ingress port and destined to a certain egress port are always delivered in the same order in which they were received.

Use IOD only if your environment cannot support out-of-order frame delivery.




Tip If you enable the in-order delivery feature, the graceful shutdown feature is not implemented.

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 61: Route Change Delivery




If the in-order guarantee feature is enabled, the frames within the network are treated as follows:

- Frames in the network are delivered in the order in which they are transmitted.
- Frames that cannot be delivered in order within the network latency drop period are dropped inside the network.

About Reordering PortChannel Frames

When a link change occurs in a PortChannel, the frames for the same exchange or the same flow can switch from one path to another faster path.

Figure 62: Link Congestion Delivery



The in-order delivery feature attempts to minimize the number of frames dropped during PortChannel link changes when the in-order delivery is enabled by sending a request to the remote switch on the PortChannel to flush all frames for this PortChannel.

**Note**

Both switches on the PortChannel must be running Cisco SAN-OS Release 3.0(1) for this IOD enhancement. For earlier releases, IOD waits for the switch latency period before sending new frames.

When the in-order delivery guarantee feature is enabled and a PortChannel link change occurs, the frames crossing the PortChannel are treated as follows:

- Frames using the old path are delivered before new frames are accepted.
- The new frames are delivered through the new path after the switch latency drop period has elapsed and all old frames are flushed.

Frames that cannot be delivered in order through the old path within the switch latency drop period are dropped. See the [Configuring the Drop Latency Time](#), on page 502.

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 MDS 9000 Family.

**Tip**

We recommend that you only enable this feature when devices that cannot handle any out-of-order frames are present in the switch. Load-balancing algorithms within the Cisco MDS 9000 Family ensure that frames are delivered in order during normal fabric operation. The load-balancing algorithms based on source FC ID, destination FC ID, and exchange ID are enforced in hardware without any performance degradation. However, if the fabric encounters a failure and this feature is enabled, the recovery will be delayed because of an intentional pausing of fabric forwarding to purge the fabric of resident frames that could potentially be forwarded out-of-order.

About Flow Statistics

If you enable flow counters, you can enable a maximum of 1 K entries for aggregate flow and flow statistics for Generation 1 modules, and 2 K entries for Generation 2 modules. 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.

Generation 1 modules allow a maximum of 1024 flow statements per module. Generation 2 modules allow a maximum of 2048-128 flow statements per module.



Note For each session, fcfow counter will increment only on locally connected devices and should be configured on the switch where the initiator is connected.

Licensing Requirements for FSPF

The following table shows the licensing requirements for this feature:

License	License Description
ENTERPRISE_PKG	The enterprise license is required to enable Fibre Channel routing services and protocols. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .

Default Settings

[Table 59: Default FSPF Settings](#), on page 497 lists the default settings for FSPF features.

Table 59: 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.

Parameters	Default
Drop latency	Disabled.
Static route cost	If the cost (metric) of the route is not specified, the default is 10.
Remote destination switch	If the remote destination switch is not specified, the default is direct.
Multicast routing	Uses the principal switch to compute the multicast tree.

Configuring FSPF

This section includes the following topics:

Configuring FSPF on a VSAN

To configure an FSPF feature for the entire VSAN, follow these steps:

Procedure

-
- Step 1** Expand a Fabric, expand a VSAN and select FSPF for a VSAN that you want to configure for FSPF.
You see the FSPF configuration in the Information pane.
 - Step 2** The RegionID, **Spf Comp Holdtime**, LSR Min Arrival, and LSR Min Interval field values are applied across all interfaces on the VSAN. You can change them here or, if they do not exist create them here.
 - Step 3** Click Apply Changes to save these changes, or click Undo Changes to discard any unsaved changes.
-

Resetting FSPF to the Default Configuration

To return the FSPF VSAN global configuration to its factory default, follow these steps:

Procedure

-
- Step 1** Expand a Fabric, expand a VSAN, and select FSPF for a VSAN that you want to configure for FSPF.
You see the FSPF configuration in the Information pane.
 - 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, follow these steps:

Procedure

- Step 1** Expand a Fabric, expand a VSAN, and select FSPF for a VSAN that you want to configure for FSPF.
You see the FSPF configuration in the Information pane.
- Step 2** Set the Status Admin drop-down menu 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.
-

Configuring FSPF Link Cost

To configure FSPF link cost, follow these steps:

Procedure

- Step 1** Expand Switches, expand FC Interfaces, and then select 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.
- 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.
-

Configuring Hello Time Intervals

To configure the FSPF Hello time interval, follow these steps:

Procedure

- Step 1** Expand Switches, expand FC Interfaces, and then select 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.
- 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.
-

Configuring Dead Time Intervals

To configure the FSPF dead time interval, follow these steps:

Procedure

-
- Step 1** Expand Switches, expand FC Interfaces, and then select 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.
- 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.
-

Configuring Retransmitting Intervals

To configure the FSPF retransmit time interval, follow these steps:

Procedure

-
- Step 1** Expand Switches, expand FC Interfaces, and then select 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.
- 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.
-

Disabling FSPF for Specific Interfaces

To disable FSPF for a specific interface, follow these steps:

Procedure

-
- Step 1** Expand Switches, expand FC Interfaces, and then select 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.
- Step 3** Set a switch Admin Status drop-down menu to down.
- Step 4** Click Apply Changes to save these changes, or click Undo Changes to discard any unsaved changes.
-

Configuring Fibre Channel Routes

To configure a Fibre Channel route using Device Manager, follow these steps:

Procedure

- Step 1** Click FC > Advanced > Routes.
You see the FC Static Route Configuration dialog box.
 - Step 2** Click Create to create a static route.
You see the Create Route dialog box.
 - Step 3** Select the VSAN ID that 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** Select the interface that you want to use to reach this destination.
 - Step 6** Select the next hop domain ID and route metric.
 - Step 7** Select either the local or remote radio button.
 - Step 8** Click Create to save these changes or click Close to discard any unsaved changes.
-

Setting the Multicast Root Switch

To use the lowest domain switch for the multicast tree computation, follow these steps:

Procedure

- Step 1** Expand a fabric, expand a VSAN, and then select Advanced for the VSAN that you want to configure FSPF on.
You see the advanced Fibre Channel configuration in the Information pane.
 - Step 2** Click the Multicast Root tab.
You see the multicast root configuration in the Information pane.
 - Step 3** Set the Config Mode drop-down menu to lowestDomainSwitch.
 - Step 4** Click Apply Changes to save these changes or click Undo Changes to discard any unsaved changes.
-

Enabling In-Order Delivery for a VSAN

To use the lowest domain switch for the multicast tree computation, follow these steps:

Procedure

- Step 1** Expand a fabric and select **All VSANS**.

- Step 2** Select the Attributes tab.
You see the general VSAN attributes in the Information pane.
- 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

To configure the drop latency time for a switch, follow these steps:

Procedure

- Step 1** Expand a fabric and select All VSANS.
You see the VSAN configuration in the Information pane.
- Step 2** Click the Attributes tab.
You see the general VSAN attributes in the Information pane.
- 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.

Verifying FSPF Configuration

To display FSPF configuration information, perform one of the following tasks:

Command	Purpose
show in-order-guarantee	Displays the In-Order Delivery Status
show fcdroplateny	Displays Latency information
show fcf flow stats aggregated module 2	Displays Aggregated Flow Details for the Specified Module
show fcf flow stats module 2	Displays Flow Details for the Specified Module
show fcf flow stats usage module 2	Displays Flow Index Usage for the Specified Module
show fspf vsan 1	Displays FSPF Information for a Specified VSAN
show fspf database vsan 1	Displays FSPF Database Information
show fspf vsan 1 interface fc1/1	Displays FSPF Interface Information

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference* .

This section contains the following topics:

Displaying the FSPF Database

The FSPF database for a specified VSAN includes the following information:

- Link State Record (LSR) type
- Domain ID of the LSR owner
- Domain ID of the advertising router
- LSR age
- LSR incarnation member
- Number of links

To display the FSPF database using Device Manager, follow these steps:

Procedure

- Step 1** Choose FC > Advanced > FSPF.
You see the FSPF dialog box.
- Step 2** Click the LSDB LSRs tab.
You see the FSPF database information.
- Step 3** Click Close to close the dialog box.
-

Displaying FSPF Statistics

To view FSPF statistics using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand a Fabric, expand a VSAN, and then select FSPF in the Logical Domains pane.
You see the FSPF configuration dialog box.
- Step 2** Click the Statistics tab.
You see the FSPF VSAN statistics in the Information pane.
- Step 3** Click the Interface Statistics tab.
You see the FSPF interface statistics in the Information pane.
-

Configuration Examples for FSPF

This section provides examples of topologies and applications that demonstrate the benefits of FSPF.




Note The FSPF feature can be used on any topology.

Fault Tolerant Fabric

[Figure 63: Fault Tolerant Fabric , on page 504](#) 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 63: Fault Tolerant Fabric




For example, if all links are of equal speed, the FSPF calculates two equal paths from A to C: A-D-C (green) and A-E-C (blue).

Redundant Links

To further improve on the topology in [Figure 63: Fault Tolerant Fabric , on page 504](#), each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. [Figure 64: Fault Tolerant Fabric with Redundant Links , on page 504](#) shows this arrangement. Because switches in the Cisco MDS 9000 Family support PortChanneling, each pair of physical links can appear to the FSPF protocol as one single logical link.

By bundling pairs of physical links, FSPF efficiency is considerably improved by the reduced database size and the frequency of link updates. Once physical links are aggregated, failures are not attached to a single link but to the entire PortChannel. This configuration also improves the resiliency of the network. The failure of a link in a PortChannel does not trigger a route change, thereby reducing the risks of routing loops, traffic loss, or fabric downtime for route reconfiguration.

Figure 64: Fault Tolerant Fabric with Redundant Links



For example, if all links are of equal speed and no PortChannels exist, the FSPF calculates four equal paths from A to C: A1-E-C, A2-E-C, A3-D-C, and A4-D-C. If PortChannels exist, these paths are reduced to two.

Failover Scenarios for PortChannels and FSPF Links

The SmartBits traffic generator was used to evaluate the scenarios displayed in [Figure 65: Failover Scenario Using Traffic Generators, on page 505](#). Two links between switch 1 and switch 2 exist as either equal-cost

ISLs or PortChannels. There is one flow from traffic generator 1 to traffic generator 2. The traffic was tested at 100 percent utilization at 1 Gbps in two scenarios:

- Disabling the traffic link by physically removing the cable (see [Table 60: Physically Removing the Cable for the SmartBits Scenario, on page 505](#)).
- Shutting down either switch 1 or switch 2 (see [Table 61: Shutting Down the Switch for the SmartBits Scenario, on page 505](#)).

Figure 65: Failover Scenario Using Traffic Generators




Table 60: Physically Removing the Cable for the SmartBits Scenario

PortChannel Scenario		FSPF Scenario (Equal cost ISL)	
Switch 1	Switch 2	Switch 1	Switch 2
110 msec (~2K frame drops)		130+ msec (~4k frame drops)	
100 msec (hold time when a signal loss is reported as mandated by the standard)			

Table 61: Shutting Down the Switch for the SmartBits Scenario

PortChannel Scenario		FSPF Scenario (Equal cost ISL)	
Switch 1	Switch 2	Switch 1	Switch 2
~0 msec (~8 frame drops)	110 msec (~2K frame drops)	130+ msec (~4K frame drops)	
No hold time needed	Signal loss on switch 1	No hold time needed	Signal loss on switch 1

Field Descriptions for FSPF

This section displays the field descriptions for this feature.

FSPF General

Field	Description
AdminStatus	The desired state of FSPF on this VSAN.
OperStatus	State of FSPF on this VSAN.

Field	Description
SetToDefault	Enabling this changes each value in this row to its default value. If all the configuration parameters have their default values and if the VSAN is suspended, then the row is deleted automatically.
RegionId	The autonomous region of the local switch on this VSAN.
DomainId	The domain ID of the local switch on this VSAN.
SpfHoldTime	The minimum time between two consecutive SPF computations on this VSAN. The smaller value means that routing will react to the changes faster but the CPU usage is greater.
SpfDelay	The time between when FSPF receives topology updates and when it starts the Shortest Path First (SPF) computation on this VSAN. The smaller value means that routing will react to the changes faster but the CPU usage is greater.
MinLsArrival	The minimum time after accepting a Link State Record (LSR) on this VSAN before accepting another update of the same LSR on the same VSAN. An LSR update that is not accepted because of this time interval is discarded.
MinLsInterval	The minimum time after this switch sends an LSR on this VSAN before it will send another update of the same LSR on the same VSAN.
LsRefreshTime	The interval between transmission of refresh LSRs on this VSAN.
LSRMaxAge	The maximum age an LSR will be retained in the FSPF database on this VSAN. It is removed from the database after MaxAge is reached.
CreateTime	When this entry was last created.
Checksum	The total checksum of all the LSRs on this VSAN.

FSPF Interfaces

Field	Description
SetToDefault	Enabling this changes each value in this row to its default value. If all the configuration parameters have their default values and if the interface is down, then the row is deleted automatically.
Cost	The administrative cost of sending a frame on this interface on this VSAN. The value 0 means that the cost has not been configured. Once the value has been configured, the value cannot again be 0; so, obviously the value cannot be set to 0. If the value is 0 and the corresponding interface is up, the agent sets a value calculated using the ifSpeed of the interface. Otherwise, the value is used as the cost. The following formula is used to calculate the link cost: $\text{Link Cost} = \begin{cases} \text{fspflfCost} & \text{if } \text{fspflfCost} > 0 \\ (1.0625 \times 10^{12} / \text{Baud Rate}) & \text{if } \text{fspflfCost} == 0 \end{cases}$ where Baud Rate is the ifSpeed of the interface.
AdminStatus	The desired state of FSPF on this interface on this VSAN.

Field	Description
HelloInterval	Interval between the periodic hello messages sent on this interface on this VSAN to verify the link health. Note that this value must be same on both the interfaces on each end of the link on this VSAN.
DeadInterval	Maximum time for which no hello messages can be received on this interface on this VSAN. After this time, the interface is assumed to be broken and removed from the database. This value must be greater than the hello interval specified on this interface on this VSAN.
RetransmitInterval	Time after which an unacknowledged link update is retransmitted on this interface on this VSAN.
Neighbour State	The state of FSPF's neighbor state machine, which is the operational state of the interaction with the neighbor's interface which is connected to this interface.
Neighbour DomainId	The domain ID of the neighbor on this VSAN.
Neighbour PortIndex	The index, as known by the neighbor, of the neighbor's interface which is connected to this interface on this VSAN.
CreateTime	When this entry was last created.

FSPF Interface Stats

Field	Description
CreateTime	When this entry was last created.
ErrorRxPkts	Number of invalid FSPF control frames received on this interface on this VSAN since the creation of the entry.
InactivityExpirations	Number of times the inactivity timer has expired on this interface on this VSAN since the creation of the entry.
LsuRxPkts	Number of Link State Update (LSU) frames received on this interface on this VSAN since the creation of the entry.
LsuTxPkts	Number of Link State Update (LSU) frames transmitted on this interface on this VSAN since the creation of the entry.
RetransmittedLsuTxPkts	Number of LSU frames retransmitted on this interface on this VSAN since the creation of the entry.
LsaRxPkts	Number of Link State Acknowledgement (LSA) frames received on this interface on this VSAN since the creation of the entry.
LsaTxPkts	Number of Link State Acknowledgement (LSA) frames transmitted on this interface on this VSAN since the creation of the entry.
HelloTxPkts	Number of HELLO frames transmitted on this interface on this VSAN since the creation of the entry.

Field	Description
HelloRxPkts	Number of HELLO frames received on this interface on this VSAN since the creation of the entry.

FSPF LSDB Links

Field	Description
NbrDomainId	The domain ID of the neighbor on the other end of this link on this VSAN.
PortIndex	The source E_port of this link, as indicated by the index value in the LSR received from the switch identified by the domain ID.
NbrPortIndex	The destination E_port of this link, as indicated by the index value in the LSR received from the switch identified by NbrDomainId.
Cost	The cost of sending a frame on this link on this VSAN. Link cost is calculated using a formula $\text{link cost} = S * (1.0625e12 / \text{Baud Rate})$ where S (value of Cost on the interface on the switch corresponding to the domain Id) is the administratively set cost factor for this interface.

FSPF LSDB LSRs

Field	Description
AdvDomainId	Domain ID of the switch that is advertising the LSR on the behalf of the switch owning it.
Age	Time since this LSR was inserted into the database.
IncarnationNumber	The link state incarnation number of this LSR. This is used to identify most recent instance of an LSR while updating the topology database when an LSR is received. The updating of an LSR includes incrementing its incarnation number prior to transmission of the updated LSR. So most recent LSR is the one with larger incarnation number.
Checksum	The checksum of the LSR.
Links	Number of entries associated with this LSR.
External	Indicates if this is an external LSR advertised by local switch.

FSPF Statistics

Field	Description
SpfComputations	The number of times the SPF computation has been done on this VSAN since the creation of the entry.
ErrorRxPkts	Number of invalid FSPF control frames received on all the interface on this VSAN since the creation of the entry.

Field	Description
ChecksumErrors	The number of FSPF checksum errors occurred on this on this VSAN since the creation of the entry.
LsuRxPkts	Total number of Link State Update (LSU) frames received on all the interfaces on this VSAN since the creation of the entry.
LsuTxPkts	Total number of Link State Update (LSU) frames transmitted on all the interfaces on this VSAN since the creation of the entry.
RetransmittedLsuTxPkts	Total number of LSU frames retransmitted on all the interfaces on this VSAN since the creation of the entry.
LsaRxPkts	Total number of Link State Acknowledgement (LSA) frames received on all the interfaces on this VSAN since the creation of the entry.
LsaTxPkts	Total number of Link State Acknowledgement (LSA) frames transmitted on all the interfaces on this VSAN since the creation of the entry.
HelloTxPkts	Total number of HELLO frames transmitted on all interfaces on this VSAN since the creation of the entry.
HelloRxPkts	Total number of HELLO frames received on all the interfaces on this VSAN since the creation of the entry.
MaxAgeCount	The number of times any LSR reached fspfMaxAge in this VSAN since the creation of the entry.

Additional References

For additional information related to implementing VSANs, see the following section:

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	Cisco MDS 9000 Family Command Reference

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none">• CISCO-FC-ROUTE-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html



CHAPTER 23

Managing FLOGI, Name Server, FDMI, and RSCN Databases

- [Managing FLOGI, Name Server, FDMI, and RSCN Databases, on page 511](#)

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, and Registered State Change Notification (RSCN) information provided in the Cisco MDS 9000 Family Switches. It includes the following sections:

Information About FLOGI

In a Fibre Channel fabric, each host or disk requires an FC ID. If the required device is displayed in the FLOGI table, the fabric login is successful. Examine the FLOGI database on a switch that is directly connected to the host HBA and connected ports. See the *Default Company ID List* section and the *Switch Interoperability* section.

In a Fibre Channel fabric, each host or disk requires an FC ID. Use the **show flogi** command to verify if a storage device is displayed in the FLOGI table as in the next section. 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.

This section includes the following topics:

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 want to modify (update or delete) the contents of a database entry that was previously registered by a different device.

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.

About Rejecting Duplicate pWWN

You can prevent a malicious or accidental login when using another device's pWWN by enabling the **reject-duplicate-pwwn** option. If you disable this option, these pWWNs are allowed to log in to the fabric and replace the first device in the name server database.

You can prevent a malicious or accidental login when 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.

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.

FDMI

Cisco MDS 9000 Family switches provide support for the Fabric-Device Management Interface (FDMI) functionality, as described in the FC-GS-4 standard. FDMI enables management of devices such as Fibre Channel host bus adapters (HBAs) through in-band communications. This addition complements the existing Fibre Channel name server and management server functions.

Using the FDMI functionality, the NX-OS software can extract the following management information about attached HBAs and host operating systems without installing proprietary host agents:

- Manufacturer, model, and serial number
- Node name and node symbolic name
- Hardware, driver, and firmware versions
- Host operating system (OS) name and version number

All FDMI entries are stored in persistent storage and are retrieved when the FDMI process is started.

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.

Apart from sending these events to registered hosts, a switch RSCN (SW-RSCN) is sent to all reachable switches in the fabric.



Note The switch sends an RSCN to notify registered nodes that a change has occurred. It is up to the nodes to query the name server again to obtain the new information. The details of the changed information are not delivered by the switch in the RSCN sent to the nodes.

About the multi-pid Option

If the RSCN **multi-pid** option is enabled, then RSCNs generated to the registered Nx ports may contain more than one affected port IDs. In this case, zoning rules are applied before putting the multiple affected port IDs together in a single RSCN. By enabling this option, you can reduce the number of RSCNs. For example, suppose you have two disks (D1, D2) and a host (H) connected to switch 1. Host H is registered to receive RSCNs. D1, D2, and H belong to the same zone. If disks D1 and D2 are online at the same time, then one of the following applies:

- The **multi-pid** option is disabled on switch 1— Two RSCNs are generated to host H, one for the disk D1 and another for disk D2.
- The **multi-pid** 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).

Some Nx ports might not support multi-pid RSCN payloads. If this situation occurs, disable the RSCN **multi-pid** option.

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.



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 Cisco MDS NX-OS releases during an upgrade or downgrade is supported by **conf-check** provided by CFS. If you attempt to downgrade from Cisco MDS SAN-OS Release 3.0, you are prompted with a **conf-check** warning. You are required to disable RSCN timer distribution support before you downgrade.

By default, the RSCN timer distribution capability is disabled and is therefore compatible when upgrading from any Cisco MDS SAN-OS release earlier than Release 3.0.

RSCN Timer Configuration Distribution

Because the timeout value for each switch is configured manually, a misconfiguration occurs when different switches time out at different times. This means different Nports in a network can receive RSCNs at different times. Cisco Fabric Services (CFS) infrastructure alleviates this situation by automatically distributing the RSCN timer configuration information to all switches in a fabric. This action also reduces the number of SW-RSCNs. Refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

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.



Note Only the RSCN timer configuration is distributed.

The RSCN timer is registered with CFS during initialization and switchover. For high availability, if the RSCN timer distribution crashes and restarts or a switchover occurs, it resumes normal functionality from the state prior to the crash or switchover.



Note You can determine the compatibility when downgrading to an earlier Cisco MDS NX-OS release using **show incompatibility system** command. You must disable RSCN timer distribution support before downgrading to an earlier release.



Note By default, the RSCN timer distribution capability is disabled and is compatible when upgrading from any Cisco MDS SAN-OS release earlier than 3.0.



Note For CFS distribution to operate correctly for the RSCN timer configuration, all switches in the fabric must be running Cisco SAN-OS Release 3.0(1) or later, or Cisco NX-OS 4.1(1b).

Locking the Fabric

The first action that modifies the database 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 along with the first active change.

Default Settings

Table 62: Default RSCN Settings , on page 515 lists the default settings for RSCN.

Table 62: Default RSCN Settings

Parameters	Default
RSCN timer value	2000 milliseconds for Fibre Channel VSANs1000 milliseconds for FICON VSANs
RSCN timer configuration distribution	Disabled

Registering Name Server Proxies

This section includes the following topics:

Registering Name Server Proxies

To register the name server proxy, follow these steps:

Procedure

-
- Step 1** Expand a fabric, expand a VSAN, and then select Advanced.
You see the VSAN advanced configuration in the Information pane.
 - Step 2** Click the NS Proxies tab.
You see the existing name server proxy for the selected VSAN.
 - Step 3** Double-click the PortName field 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.
-

Configuring the multi-pid Option

To configure the **multi-pid** option, follow these steps:

Procedure

-
- Step 1** Expand a fabric, expand a VSAN, and then select Advanced.
You see the VSAN advanced configuration in the Information pane.
 - Step 2** Click the RSCN Multi-PID tab.
 - 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.
-

Suppressing Domain Format SW-RSCNs

A domain format SW-RSCN is sent whenever the local switch name or the local switch management IP address changes. This SW-RSCN is sent to all other domains and switches over the ISLs. The remote switches can issue GMAL and GIELN commands to the switch that initiated the domain format SW-RSCN to determine what changed. Domain format SW-RSCNs can cause problems with some non-Cisco MDS switches (refer to the [Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide](#)).



Note You cannot suppress transmission of port address or area address format RSCNs.

Configuring the RSCN Timer with CFS

To configure the RSCN timer with CFS, follow these steps:

Procedure

- Step 1** Expand a fabric, expand a VSAN, and then select Advanced in the Logical Domains pane.
- Step 2** Click the RSCN Event tab.
You see the VSAN advanced configuration in the Information pane.
- 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.

Configuring the RSCN Timer

RSCN maintains a per-VSAN event list queue, where the RSCN events are queued as they are generated. When the first RSCN event is queued, a per VSAN timer starts. Upon time-out, all the events are dequeued and coalesced RSCNs are sent to registered users. The default timer values minimize the number of coalesced RSCNs sent to registered users. Some deployments require smaller event timer values to track changes in the fabric.



Note The RSCN timer value must be the same on all switches in the VSAN. See the [RSCN Timer Configuration Distribution Using CFS, on page 513](#).



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.

Committing the RSCN Timer Configuration Changes

If you commit the changes made to the active database, the configuration is committed to all the switches in the fabric. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

Discarding the RSCN Timer Configuration Changes

If you discard (abort) the changes made to the pending database, the configuration database remains unaffected and the lock is released.

Clearing a Locked Session

If you have changed the RSCN timer configuration and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip The pending database is only available in the volatile directory and are subject to being discarded if the switch is restarted.

Displaying FLOGI Details

To verify that a storage device is in the fabric login (FLOGI) table, follow these steps:

Procedure

-
- Step 1** Expand Switches, expand Interfaces, and then select FC Physical.
You see the interface configuration in the Information pane.
- Step 2** Click the FLOGI tab.
You see all end devices that are logged into the fabric.
-

Viewing Name Server Database Entries

To view the name server database using Device Manager, follow these steps:

Procedure

-
- Step 1** Select FC > Name Server.
You see the 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.

Displaying RSCN Information

To display RSCN information, follow these steps:

Procedure

- Step 1** Expand a fabric, expand a VSAN, and then select Advanced.
You see the VSAN advanced configuration in the Information pane.
- Step 2** Click the RSCN Reg tab or the RSCN Statistics tab.

Field Descriptions for Databases

This setion contains the field descriptions for this feature.

FC Interfaces FLOGI

Field	Description
FcId	The address identifier that has been assigned to the logged-in Nx_Port.
PortName	The world wide name of the logged-in Nx_Port.
NodeName	The world wide name of the Remote Node the logged-in Nx_Port belongs to.
Original PWWN	The original port WWN for this interface.
Version	The version of FC-PH that the Fx_Port has agreed to support from the Fabric Login.
BBCredit Rx	The maximum number of receive buffers available for holding Class 2, Class 3 received from the logged-in Nx_Port. It is for buffer-to-buffer flow control in the incoming direction from the logged-in Nx_Port to FC-port.
BBCredit Tx	The total number of buffers available for holding Class 2, Class 3 frames to be transmitted to the logged-in Nx_Port. It is for buffer-to-buffer flow control in the direction from FC-Port to Nx_Port. The buffer-to-buffer flow control mechanism is indicated in the respective BbCreditModel.
CoS	The classes of services that the logged-in Nx_Port has requested the FC-Port to support and the FC-Port has granted the request.
Class2 RxDataSize	The Class 2 Receive Data Field Size of the logged-in Nx_Port. Specifies the largest Data Field Size for an FT_1 frame that can be received by the Nx_Port.

Field	Description
Class2 SeqDeliv	Whether the FC-Port has agreed to support Class 2 sequential delivery during the Fabric Login. This is meaningful only if Class 2 service has been agreed. This is applicable only to Fx_Ports.
Class3 RxDataSize	The Class3 Receive Data Field Size of the logged-in Nx_Port. Specifies the largest Data Field Size for an FT_1 frame that can be received by the Nx_Port.
Class3 SeqDeliv	Whether the FxPort has agreed to support Class 3 sequential delivery during the Fabric Login. This is meaningful only if Class 3 service has been agreed. This is applicable only to Fx_Ports.

FDMI HBAs

Field	Description
Sn	The serial number of this HBA.
Model	The model of this HBA.
ModelDescr	The model description.
OSInfo	The type and version of the operating system controlling this HBA.
MaxCTPayload	The maximum size of the Common Transport (CT) payload including all CT headers but no FC frame header(s), that may be send or received by application software resident in the host containing this HBA.

FDMI Ports

Field	Description
SupportedFC4Type	The supported FC-4 types attribute registered for this port on this VSAN.
SupportedSpeed	The supported speed registered for this port on this VSAN.
CurrentSpeed	The current speed registered for this port on this VSAN.
MaxFrameSize	The maximum frame size attribute registered for this port on this VSAN.
OsDevName	The OS Device Name attribute registered for this port on this VSAN.
HostName	The name of the host associated with this port.

FDMI Versions

Field	Description
Hardware	The hardware version of this HBA.
DriverVer	The version level of the driver software controlling this HBA.

Field	Description
OptROMVer	The version of the Option ROM or the BIOS of this HBA.
Firmware	The version of the firmware executed by this HBA.

RSCN Nx Registrations

Field	Description
RegType	Indicates the type of registration desired by the subscriber. <ul style="list-style-type: none"> • 'fromFabricCtrlr' indicates RSCNs generated by the Fabric Controller. • 'fromNxPort' indicates RSCNs generated by Nx_Ports. • 'fromBoth' indicates RSCNs generated by Fabric Controller and Nx_Ports.

RSCN Multi-PID Support

Field	Description
Enable	Specifies whether the multi-pid option is enabled on this VSAN.

RSCN Event

Field	Description
TimeOut (msec)	The time (in seconds) before the RSCN event times out.

RSCN Statistics

Field	Description
SCR Rx	The number of SCRs received from Nx_Ports on this VSAN.
SCR RJT	The number of SCR rejected on this VSAN.
RSCN Rx	The number of RSCNs from Nx_Ports received on this VSAN.
RSCN Tx	The total number of RSCNs transmitted on this VSAN.
RSCN RJT	The number of RSCN requests rejected on this VSAN.
SW-RSCN Rx	The number of Inter-Switch Registered State Change Notifications (SW_RSCN) received on this VSAN from other switches.
SW-RSCN Tx	The number of Inter-Switch Registered State Change Notifications (SW_RSCN) transmitted on this VSAN to other switches.

Field	Description
SW-RSCN RJT	The number of SW_RSCN requests rejected on this VSAN.

Name Server General

Field	Description
Type	The port type of this port.
PortName	The fibre channel Port_Name (WWN) of this Nx_port.
NodeName	The fibre channel Node_Name (WWN) of this Nx_port.
FC4Type/Features	The FC-4 Features associated with this port and the FC-4 Type. Refer to FC-GS3 specification for the format.
SymbolicPortName	The user-defined name of this port.
SymbolicNodeName	The user-defined name of the node of this port.
FabricPortName	The fabric port name (WWN) of the Fx_port to which this Nx_port is attached.

Name Server Advanced

Field	Description
ClassOfSvc	The class of service indicator.
PortIpAddress	Contains the IP address of the associated port.
NodeIpAddress	The IP address of the node of this Nx_port, as indicated by the Nx_Port in a GS3 message that it transmitted.
SymbolicPortName	The user-defined name of this port.
SymbolicNodeName	The user-defined name of the node of this port.
HardAddress	Extended Link Service (FC-PH-2). Hard Address is the 24-bit NL_Port identifier which consists of - the 8-bit Domain Id in the most significant byte - the 8-bit Area Id in the next most significant byte - the 8-bit AL-PA(Arbitrated Loop Physical Address) which an NL_port attempts acquire during FC-AL initialization in the least significant byte. If the port is not an NL_Port, or if it is an NL_Port but does not have a hard address, then all bits are reported as 0s.
ProcAssoc	The Fibre Channel initial process associator (IPA).
PermanentPortName	The permanent port name of this Nx port. If multiple port names are associated with this Nx port via FDISC (Discover F Port Service Parameters), the permanent port name is the original port name associated with this Nx port at login.

Name Server Proxy

Field	Description
PortName	Name of the proxy port which can register or deregister for other ports on this VSAN. Users can enable third-party registrations by setting this value.

Name Server Statistics

Field	Description
Queries Rx	The total number of Get Requests received by the local switch on this VSAN.
Queries Tx	The total number of Get Requests sent by the local switch on this VSAN.
Requests Rx Reg	The total number of Registration Requests received by the local switch on this VSAN.
Requests Rx DeReg	The total number of De-registration Requests received by the local switch on this VSAN.
RSCN Rx	The total number of RSCN commands received by the local switch on this VSAN.
RSCN Tx	The total number of RSCN commands sent by the local switch on this VSAN.
Rejects Tx	The total number of requests rejected by the local switch on this VSAN.



CHAPTER 24

Configuring FICON

- [Configuring FICON, on page 523](#)

Configuring FICON

Fibre Connection (FICON) interface capabilities enhance the Cisco MDS 9000 Family by supporting both open systems and mainframe storage network environments. The control unit port (CUP) also is supported, which allows in-band management of the switch from FICON processors.



Note

Cisco Fabric Manager release 3.x does not support FICON management of Cisco MDS 9000 Family switches running SAN-OS release 2.(x).

Information About FICON

The Cisco MDS 9000 Family supports the Fibre Channel Protocol (FCP), FICON, iSCSI, and FCIP capabilities within a single, high-availability platform (see [Figure 66: Shared System Storage Network, on page 524](#)).


The FICON feature is not supported on:

- Cisco MDS 9120 switches
- Cisco MDS 9124 switches
- Cisco MDS 9140 switches
- The 32-port Fibre Channel switching module
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeSystem

FCP and FICON are different FC4 protocols and their traffic is independent of each other. Devices using these protocols should be isolated using VSANs.

The fabric binding feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations (refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*). The Registered Link Incident Report (RLIR) application provides a method for a switch port to send an LIR to a registered Nx port.

Figure 66: Shared System Storage Network



This section includes the following topics:

FICON Requirements

The FICON feature has the following requirements:

- You can implement FICON features in the following switches:
 - Any switch in the Cisco MDS 9500 Series
 - Any switch in the Cisco MDS 9200 Series (including the Cisco MDS 9222i Multiservice Modular Switch)
 - Cisco MDS 9134 Multilayer Fabric Switch
 - MDS 9000 Family 18/4-Port Multiservice Module
- You need the MAINFRAME_PKG license to configure FICON parameters.
- To extend your FICON configuration over a WAN link using FCIP, you need the appropriate SAN_EXTN_OVER_IP license for the module you are using. For more information, refer to the *Cisco NX-OS Family Licensing Guide*.

Cisco MDS-Specific FICON Advantages


This section explains the additional FICON advantages in Cisco MDS switches and includes the following topics:

Fabric Optimization with VSANs

Generally, separate physical fabrics have a high level of switch management and have a higher implementation cost. The ports in each island also may be over-provisioned depending on the fabric configuration.

By using the Cisco MDS-specific VSAN technology, you can have greater efficiency between these physical fabrics by lowering the cost of over-provisioning and reducing the number of switches to be managed. VSANs also help you to move unused ports nondisruptively and provide a common redundant physical infrastructure (see [Figure 67: VSAN-Specific Fabric Optimization, on page 525](#)).

Figure 67: VSAN-Specific Fabric Optimization



VSANs enable global SAN consolidation by allowing you to convert existing SAN islands into virtual SAN islands on a single physical network. It provides hardware-enforced security and separation between applications or departments to allow coexistence on a single network. It also allows virtual rewiring to consolidate your storage infrastructure. You can move assets between departments or applications without the expense and disruption of physical relocation of equipment.



Note You can configure VSANs in any Cisco MDS switch, but you only can enable FICON in up to eight of these VSANs. The number of VSANs configured depends on the platform.

Mainframe users can think of VSANs as being like FICON LPARs in the MDS SAN fabric. You can partition switch resources into FICON LPARs (VSANs) that are isolated from each other, in much the same way that you can partition resources on a zSeries or DS8000. Each VSAN has its own set of fabric services (such as fabric server and name server), FICON CUP, domain ID, Fabric Shortest Path First (FSPF) routing, operating mode, IP address, and security profile. FICON LPARs can span line cards and are dynamic in size. For example, one FICON LPAR with 10 ports can span 10 different line cards. FICON LPARs can also include ports on more than one switch in a cascaded configuration. The consistent fairness of the Cisco MDS 9000 switching architecture means that “all ports are created equal,” simplifying provisioning by eliminating the “local switching” issues seen on other vendors’ platforms. Addition of ports to a FICON LPAR is a nondisruptive process. The maximum number of ports for a FICON LPAR is 255 due to FICON addressing limitations.

FCIP Support

The multilayer architecture of the Cisco MDS 9000 Family enables a consistent feature set over a protocol-agnostic switch fabric. Cisco MDS 9500 Series and 9200 Series switches transparently integrate Fibre Channel, FICON, and Fibre Channel over IP (FCIP) in one system. The FICON over FCIP feature enables cost-effective access to remotely located mainframe resources. With the Cisco MDS 9000 Family platform, storage replication services such as IBM PPRC and XRC can be extended over metro to global distances using ubiquitous IP infrastructure which simplifies business continuance strategies.

Refer to the *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*.

PortChannel Support

The Cisco MDS implementation of FICON provides support for efficient utilization and increased availability of Inter-Switch Links (ISLs) necessary to build stable large-scale SAN environments. PortChannels ensure an enhanced ISL availability and performance in Cisco MDS switches.

Refer to the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide* for more information on PortChannels.

VSANs for FICON and FCP Mixing

Cisco MDS 9000 Family FICON-enabled switches simplify deployment of even the most complex mixed environments. Multiple logical FICON, Z-Series Linux/FCP, and Open-Systems Fibre Channel Protocol (FCP) fabrics can be overlaid onto a single physical fabric by simply creating VSANs as required for each service. VSANs provide both hardware isolation and protocol specific fabric services, eliminating the complexity and potential instability of zone-based mixed schemes.

By default, the FICON feature is disabled in all switches in the Cisco MDS 9000 Family. When the FICON feature is disabled, FC IDs can be allocated seamlessly. Mixed environments are addressed by the Cisco NX-OS software. The challenge of mixing FCP and FICON protocols are addressed by Cisco MDS switches when implementing VSANs.

Switches and directors in the Cisco MDS 9000 Family support FCP and FICON protocol mixing at the port level. If these protocols are mixed in the same switch, you can use VSANs to isolate FCP and FICON ports.

**Tip**

When creating a mixed environment, place all FICON devices in one VSAN (other than the default VSAN) and segregate the FCP switch ports in a separate VSAN (other than the default VSAN). This isolation ensures proper communication for all connected devices.

Cisco MDS-Supported FICON Features

The Cisco MDS 9000 Family FICON features include:

- Flexibility and investment protection—The Cisco MDS 9000 Family shares common switching and service modules across the Cisco MDS 9500 Series and the 9200 Series.

Refer to the *Cisco MDS 9500 Series Hardware Installation Guide* and the *Cisco MDS 9200 Series Hardware Installation Guide*.

- High-availability FICON-enabled director—The Cisco MDS 9500 Series combines nondisruptive software upgrades, stateful process restart and failover, and full redundancy of all major components for a new standard in director-class availability. It supports up to 528 autosensing, 4/2/1-Gbps, 10-Gbps, FICON or FCP ports in any combination in a single chassis. Refer to the *Cisco MDS 9000 Family NX-OS High Availability and Redundancy Configuration Guide*.

- Infrastructure protection—Common software releases provide infrastructure protection across all Cisco MDS 9000 platforms. Refer to the *Cisco MDS 9000 Family NX-OS Software Upgrade and Downgrade Guide*.
- VSAN technology—The Cisco MDS 9000 Family provides VSAN technology for hardware-enforced, isolated environments within a single physical fabric for secure sharing of physical infrastructure and enhanced FICON mixed support. See [Chapter 19, “Configuring and Managing VSANs.”](#)
- Port-level configurations—There are BB_credits, beacon mode, and port security for each port. Refer to the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide* for information about buffer-to-buffer credits, beacon LEDs, and trunking.
- Alias name configuration—Provides user-friendly aliases instead of the WWN for switches and attached node devices. See [Chapter 16, “Configuring and Managing Zones.”](#)
- Comprehensive security framework—The Cisco MDS 9000 Family supports RADIUS and TACACS+ authentication, Simple Network Management Protocol Version 3 (SNMPv3), role-based access control, Secure Shell Protocol (SSH), Secure File Transfer Protocol (SFTP), VSANs, hardware-enforced zoning, ACLs, fabric binding, Fibre Channel Security Protocol (FC-SP), LUN zoning, read-only zones, and VSAN-based access control. Refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide* for information about RADIUS, TACACS+, FC-SP, and DHCHAP.
- Traffic encryption—IPsec is supported over FCIP. You can encrypt FICON and Fibre Channel traffic that is carried over FCIP. Refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.
- Local accounting log—View the local accounting log to locate FICON events. For more information about MSCHAP authentication, and local AAA services, refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.
- Unified storage management—Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for in-band management using the IBM S/A OS/390 I/O operations console. See the [CUP In-Band Management, on page 540](#).
- Port address-based configurations—Configure port name, blocked or unblocked state, and the prohibit connectivity attributes can be configured on the ports. See the [Configuring FICON Ports, on page 552](#).
- You can display the following information:
 - Individual Fibre Channel ports, such as the port name, port number, Fibre Channel address, operational state, type of port, and login data.
 - Nodes attached to ports.
 - Port performance and statistics.

See the [Calculating FICON Flow Load Balance, on page 558](#).

- Configuration files—Store and apply configuration files. See the [FICON Configuration Files, on page 537](#).
- FICON and Open Systems Management Server features if installed. —See the [VSANs for FICON and FCP Mixing, on page 526](#).
- Enhanced cascading support—See the [CUP In-Band Management, on page 540](#).
- Date and time—Set the date and time on the switch. See the [Allowing the Host to Control the Timestamp, on page 550](#).
- Configure SNMP trap recipients and community names—See the [Configuring SNMP Control of FICON Parameters, on page 551](#).
- Call Home configurations—Configure the director name, location, description, and contact person. Refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.
- Configure preferred domain ID, FC ID persistence, and principal switch priority—For information about configuring domain parameters, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

- Sophisticated SPAN diagnostics—The Cisco MDS 9000 Family provides industry-first intelligent diagnostics, protocol decoding, and network analysis tools as well as integrated Call Home capability for added reliability, faster problem resolution, and reduced service costs. For information about monitoring network traffic using SPAN, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.
- Configure R_A_TOV, E_D_TOV— See the [“Fibre Channel Time-Out Values” section on page 27-2](#).
- Director-level maintenance tasks—Perform maintenance tasks for the director including maintaining firmware levels, accessing the director logs, and collecting data to support failure analysis. For information about monitoring system processes and logs refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.
- Port-level incident alerts—Display and clear port-level incident alerts. See the [Clearing RLIR Information, on page 554](#).

FICON Cascading

The Cisco MDS NX-OS software allows multiple switches in a FICON network. To configure multiple switches, you must enable and configure fabric binding in that switch (see the [Calculating FICON Flow Load Balance, on page 558](#) and refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*).

FICON VSAN Prerequisites

To ensure that a FICON VSAN is operationally up, be sure to verify the following requirements:

- Set the default zone to permit, if you are not using the zoning feature. See the [“About the Default Zone” section on page 16-5](#).
- Enable in-order delivery on the VSAN. See [Chapter 22, “Configuring Fibre Channel Routing Services and Protocols.”](#)
- Enable (and if required, configure) fabric binding on the VSAN. See the [Calculating FICON Flow Load Balance, on page 558](#). For more information about Fabric Binding, refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.
- Verify that conflicting persistent FC IDs do not exist in the switch. For information about configuring domain parameters, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.
- Verify that the configured domain ID and requested domain ID match. For information about configuring domain parameters, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.
- Add the CUP (area FE) to the zone, if you are using zoning. See the [CUP In-Band Management, on page 540](#).

If any of these requirements are not met, the FICON feature cannot be enabled.

FICON Port Numbering


With reference to the FICON feature, ports in Cisco MDS switches are identified by a statically defined 8-bit value known as the port number. A maximum of 255 port numbers are available. You can use the following port numbering schemes:

- Default port numbers based on the chassis type
- Reserved port numbers

Default FICON Port Numbering Scheme

Default FICON port numbers are assigned by the Cisco MDS NX-OS software based on the module and the slot in the chassis. The first port in a switch always starts with a zero (0) (see [Figure 68: Default FICON Port Number in Numbering on the Cisco MDS 9000 Family Switch, on page 529](#)).

Figure 68: Default FICON Port Number in Numbering on the Cisco MDS 9000 Family Switch



The default FICON port number is assigned based on the front panel location of the port and is specific to the slot in which the module resides. Thirty-two (32) port numbers are assigned to each slot on all Cisco MDS 9000 Family switches except for the Cisco MDS 9513 Director, which has 16 port numbers assigned for each slot. These default numbers are assigned regardless of the module's physical presence in the chassis, the port status (up or down), or the number of ports on the module (4, 12, 16, 24, or 48). If a module has fewer ports than the number of port numbers assigned to the slot, then the excess port numbers are unused. If a module has more ports than the number of port numbers assigned to the slot, the excess ports cannot be used for FICON traffic unless you manually assign the port numbers.



Note

You can use the **ficon slot assign port-numbers** command to make use of any Follow the steps in [Assigning FICON Port Numbers to Slots, on page 542](#) to make use of excess ports by manually assigning more port numbers to the slots. Before doing this, however, we recommend that you review the default port number assignments for Cisco MDS 9000 switches shown in [Table 65: Default FICON Settings , on page 542](#) and [Table 63: Default FICON Port Numbering in the Cisco MDS 9000 Family, on page 530](#), and that you read the following sections to gain a complete understanding of FICON port numbering: [About the Reserved FICON Port Numbering Scheme, on page 532](#), [FICON Port Numbering Guidelines, on page 541](#), and [Assigning FICON Port Numbers to Slots, on page 542](#).



Note

Only Fibre Channel, PortChannel, and FCIP ports are mapped to FICON port numbers. Other types of interfaces do not have a corresponding port number.

[Table 63: Default FICON Port Numbering in the Cisco MDS 9000 Family, on page 530](#) lists the default port number assignment for the Cisco MDS 9000 Family of switches and directors.

Table 63: Default FICON Port Numbering in the Cisco MDS 9000 Family

Product	Slot Number	Implemented Port Allocation		Unimplemented Ports	Notes
		To Ports	To PortChannel/FCIP		
Cisco MDS 9200 Series	Slot 1	0 through 31	64 through 89	90 through 253 and port 255	Similar to a switching module.
	Slot 2	32 through 63			
Cisco MDS 9222i Series	Slot 1	0 through 31	64 through 89	90 through 253 and port 255	The first 4, 12, 16, or 24 port numbers in a 4-port, 12-port, 16-port, or 24-port module are used and the rest remain unused. Extra 16 ports on 48-port modules are not allocated numbers.
	Slot 2	32 through 63			
Cisco MDS 9506 Director	Slot 1	0 through 31	128 through 153	154 through 253 and port 255	Supervisor modules are not allocated port numbers.
	Slot 2	32 through 63			
	Slot 3	64 through 95			
	Slot 4	96 through 127			
	Slot 5	None			
	Slot 6	None			
Cisco MDS 9134 Director	Slot 1	0 through 33	34 through 59	60 through 253 and port 255	

Cisco MDS 9509 Director	Slot 1	0 through 31	224 through 249	250 through 253 and port 255	The first 4, 12, 16, or 24 port numbers in a 4-port, 12-port, 16-port, or 24-port module are used and the rest remain unused. Extra 16 ports on 48-port modules are not allocated port numbers.
	Slot 2	32 through 63			
	Slot 3	64 through 95			
	Slot 4	96 through 127			
	Slot 5	None			Supervisor modules are not allocated port numbers.
	Slot 6	None			
	Slot 7	128 through 159			The first 4, 12, 16, or 24 port numbers are used for a 4-port, 12-port, 16-port, or 24-port module and the rest remain unused. Extra 16 ports on 48-port modules are not allocated port numbers.
	Slot 8	160 through 191			
	Slot 9	192 through 223			

Cisco MDS 9513 Director	Slot 1	0 through 15	224 through 249	250 through 253 and port 255	The first 4, 12 or 16 port numbers are used for a 4-port, 12-port or 16-port module and the rest remain unused. Extra ports on 24-port, 32-port, and 48-port modules are not allocated port numbers.
	Slot 2	16 through 31			
	Slot 3	32 through 47			
	Slot 4	48 through 63			
	Slot 5	64 through 79			
	Slot 6	80 through 95			
	Slot 7	None			Supervisor modules are not allocated port numbers.
	Slot 8	None			
	Slot 9	96 through 111			The first 4 or 12 port numbers are used for a 4-port or 12-port module and the rest remain unused. Extra ports on 24-port, 32-port, and 48-port modules are not allocated port numbers.
	Slot 10	112 through 127			
	Slot 11	128 through 143			
	Slot 12	144 through 159			
	Slot 13	160 through 175			

Port Addresses

By default, port numbers are the same as port addresses. You can swap the port addresses (see the [Port Swapping](#), on page 538).

Implemented and Unimplemented Port Addresses

An implemented port refers to any port address that is assigned by default to a slot in the chassis (see *Default Settings*). An unimplemented port refers to any port address that is not assigned by default to a slot in the chassis (see *Default Settings*).

About the Reserved FICON Port Numbering Scheme

A range of 250 port numbers are available for you to assign to all the ports on a switch. *Default Settings* shows that you can have more than 250 physical ports on a switch and the excess ports do not have port numbers in the default numbering scheme. When you have more than 250 physical ports on your switch, you can have ports without a port number assigned if they are not in a FICON VSAN, or you can assign duplicate port numbers if they are not used in the same FICON VSAN. For example, you can configure port number 1 on interface fc1/1 in FICON VSAN 10 and fc10/1 in FICON VSAN 20.



Note A VSAN can have a maximum of 250 port numbers.



Note FICON port numbers are not changed for ports that are active. You must first disable the interfaces using the **shutdown** command.



Note You can configure port numbers even when no module is installed in the slot.

Installed and Uninstalled Ports

An installed port refers to a port for which all required hardware is present. A specified port number in a VSAN can be implemented, and yet not installed, if any of the following conditions apply:

- The module is not present—For example, if module 1 is not physically present in slot 1 in a Cisco MDS 9509 Director, ports 0 to 31 are considered uninstalled.
- The small form-factor pluggable (SFP) port is not present—For example, if a 16-port module is inserted in slot 2 in a Cisco MDS 9509 Director, ports 48 to 63 are considered uninstalled.
- For slot 1, ports 0 to 31, or 0 to 15 have been assigned. Only the physical port fc1/5 with port number 4 is in VSAN 2. The rest of the physical ports are not in VSAN 2. The port numbers 0 to 249 are considered implemented for any FICON-enabled VSAN. Therefore, VSAN 2 has port numbers 0 to 249 and one physical port, fc1/4. The corresponding physical ports 0 to 3, and 5 to 249 are not in VSAN 2. When the FICON VSAN port address is displayed, those port numbers with the physical ports not in VSAN 2 are not installed (for example, ports 0 to 3, or 5 to 249).

Another scenario is if VSANs 1 through 5 are FICON-enabled, and trunking-enabled interface fc1/1 has VSANs 3 through 10, then port address 0 is uninstalled in VSAN 1 and 2.

- The port is part of a PortChannel—For example, if interface fc 1/1 is part of PortChannel 5, port address 0 is uninstalled in all FICON VSANs. See *Default Settings*.

About Port Numbers for FCIP and PortChannel

FCIP and PortChannels cannot be used in a FICON-enabled VSAN unless they are explicitly bound to a port number.

See the [Configuring FICON Ports, on page 552](#), [Configuring FICON Ports, on page 552](#), [Reserving FICON Port Numbers for FCIP and PortChannel Interfaces, on page 543](#), and [Binding Port Numbers to FCIP Interfaces, on page 552](#).

You can use the default port numbers if they are available (see *Default FICON Port Numbering Scheme*) or if you reserve port numbers from the pool of port numbers that are not reserved for Fibre Channel interfaces (see the [FICON Port Numbering, on page 528](#) and the [About the Reserved FICON Port Numbering Scheme, on page 532](#)).

FC ID Allocation

FICON requires a predictable and static FC ID allocation scheme. When FICON is enabled, the FC ID allocated to a device is based on the port address of the port to which it is attached. The port address forms the middle byte of the fabric address. Additionally, the last byte of the fabric address should be the same for all devices in the fabric. By default, the last byte value is 0 and can be configured (see the [Assigning FC ID Last Byte](#), on page 549).




Note

You cannot configure persistent FC IDs in FICON-enabled VSANs.

Cisco MDS switches have a dynamic FC ID allocation scheme. When FICON is enabled or disabled on a VSAN, all the ports are shut down and restarted to switch from the dynamic to static FC IDs and vice versa (see [Figure 69: Static FC ID Allocation for FICON](#), on page 534).

Figure 69: Static FC ID Allocation for FICON



About Enabling FICON on a VSAN

By default FICON is disabled in all VSANs on the switch.

You can enable FICON on a per VSAN basis in one of the following ways:

- Use the automated **setup ficon** command.

See the [Setting Up a Basic FICON Configuration](#), on page 543.

- Manually address each prerequisite.

See the [Information About FICON](#), on page 523.

- Use Device Manager (refer to the Fabric Configuration Guide, Cisco DCNM for SAN).

When you enable the FICON feature in Cisco MDS switches, the following restrictions apply:

- You cannot disable in-order delivery for the FICON-enabled VSAN.
- You cannot disable fabric binding or static domain ID configurations for the FICON-enabled VSAN.
- The load balancing scheme is changed to Source ID (SID)—Destination ID (DID). You cannot change it back to SID—DID—OXID.
- The IPL configuration file is automatically created.

See the [FICON Configuration Files](#), on page 537.

FICON Information Refresh

When viewing FICON information through the Device Manager dialog boxes, you must manually refresh the display by clicking the **Refresh** button to see the latest updates. You need to take this step whether you configure FICON through the CLI or through the Device Manager.

There is no automatic refresh of FICON information. This information would be refreshed so often that it would affect performance.

About FICON Device Allegiance

FICON requires serialization of access among multiple mainframes, CLI, and SNMP sessions be maintained on Cisco MDS 9000 Family switches by controlling device allegiance for the currently executing session. Any other session is denied permission to perform configuration changes unless the required allegiance is available.



Caution

This task discards the currently executing session.

Automatically Saving the Running Configuration

Cisco MDS NX-OS provides an option to automatically save any configuration changes to the startup configuration. This ensures that the new configuration is present after a switch reboot. By default, the Active=Saved active equals saved option is automatically enabled on any FICON VSAN.

The following table displays the results of the Active = Saved option active equals saved command and the implicit copy from the running configuration to the startup configuration (copy running start) copy running-config startup-config command in various scenarios.

When the Active=Saved option active equals saved command is enabled in any FICON-enabled VSAN in the fabric, then the following apply:

- All configuration changes (FICON-specific or not) are automatically saved to persistent storage (implicit copy running start) and stored in the startup configuration.
- FICON-specific configuration changes are immediately saved to the IPL file (see the “FICON Configuration Files” section).

If the Active=Saved option active equals saved command is not enabled in any FICON-enabled VSAN in the fabric, then FICON-specific configuration changes are not saved in the IPL file and an implicit copy running startup command is not issued, you must explicitly save the running configuration to the startup configuration issue the copy running start command explicitly.

Table 64: Saving the Active FICON and Switch Configuration

Number	FICON-enabled VSAN?	active equals saved Enabled?	Implicit copy running start Issued? ³¹	Notes
1	Yes	Yes (in all FICON VSANs)	Implicit	FICON changes written to the IPL file. Non-FICON changes saved to startup configuration and persistent storage.

Number	FICON-enabled VSAN?	active equals saved Enabled?	Implicit copy running start Issued? ³¹	Notes
2	Yes	Yes (even in one FICON VSAN)	Implicit	FICON changes written to IPL file for only the VSAN that has active equals saved option enabled. Non-FICON changes saved to startup configuration and persistent storage.
3	Yes	Not in any FICON VSAN	Not implicit	FICON changes are not written to the IPL file. Non-FICON changes are saved in persistent storage—only if you explicitly issue the copy running start command.
4	No	Not applicable		

³¹ When the Cisco NX-OS software implicitly issues a **copy running-config startup-config** command in the Cisco MDS switch, only a binary configuration is generated—an ASCII configuration is not generated). If you wish to generate an additional ASCII configuration at this stage, you must explicitly issue the **copy running-config startup-config** command again.



Note If **active equals saved** is enabled, the Cisco NX-OS software ensures that you do not have to perform the **copy running startup** command for the FICON configuration as well. If your switch or fabric consists of multiple FICON-enabled VSANs, and one of these VSANs have **active equals saved** enabled, changes made to the non-FICON configuration results in all configurations being saved to the startup configuration.

Port Prohibiting

To prevent implemented ports from talking to each other, configure prohibits between two or more ports. If you prohibit ports, the specified ports are prevented from communicating with each other.



Tip You cannot prohibit a PortChannel or FCIP interface.

Unimplemented ports are always prohibited. In addition, prohibit configurations are always symmetrically applied—if you prohibit port 0 from talking to port 15, port 15 is automatically prohibited from talking to port 0.



Note If an interface is already configured in E or TE mode and you try to prohibit that port, your prohibit configuration is rejected. Similarly, if a port is not up and you prohibit that port, the port is not allowed to come up in E mode or in TE mode.

About RLIR

The Registered Link Incident Report (RLIR) application provides a method for a switch port to send a Link Incident Record (LIR) to a registered Nx port.

When an LIR is detected in FICON-enabled switches in the Cisco MDS 9000 Family from an RLIR Extended Link Service (ELS), the switch sends that record to the members in its Established Registration List (ERL).

In case of multiswitch topology, a Distribute Registered Link Incident Record (DRLIR) Inter-Link Service (ILS) is sent to all reachable remote domains along with the RLIR ELS. On receiving the DRLIR ILS, the switch extracts the RLIR ELS and sends it to the members of the ERL.

The Nx ports interested in receiving the RLIR ELS send the Link Incident Record Registration (LIRR) ELS request to the management server on the switch. The RLIRs are processed on a per-VSAN basis.

The RLIR data is written to persistent storage when you **copy** the running configuration to the startup configuration.

FICON Configuration Files

You can save up to 16 FICON configuration files on each FICON-enabled VSAN (in persistent storage). The file format is proprietary to IBM. These files can be read and written by IBM hosts using the in-band CUP protocol. Additionally, you can use the Cisco MDS CLI or DCNM-SAN applications to operate on these FICON configuration files.



Note Multiple FICON configuration files with the same name can exist in the same switch, provided they reside in different VSANs. For example, you can create a configuration file named XYZ in both VSAN 1 and VSAN 3.

When you enable the FICON feature in a VSAN, the switches always use the startup FICON configuration file, called IPL. This file is created with a default configuration as soon as FICON is enabled in a VSAN.



Caution When FICON is disabled on a VSAN, all the FICON configuration files are irretrievably lost.

FICON configuration files contain the following configuration for each implemented port address:

- Block
- Prohibit mask
- Port address name



Note Normal configuration files used by Cisco MDS switches include FICON-enabled attributes for a VSAN, port number mapping for PortChannels and FCIP interfaces, port number to port address mapping, port and trunk allowed VSAN configuration for ports, in-order guarantee, static domain ID configuration, and fabric binding configuration.

Refer to the *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide* for details on the normal configuration files used by Cisco MDS switches.

Only one user can access the configuration file at any given time:

- If this file is being accessed by user 1, user 2 cannot access this file.
- If user 2 does attempt to access this file, an error is issued to user 2.

- If user 1 is inactive for more than 15 seconds, the file is automatically closed and available for use by any other permitted user.

FICON configuration files can be accessed by any host, SNMP, or CLI user who is permitted to access the switch. The locking mechanism in the Cisco NX-OS software restricts access to one user at a time per file. This lock applies to newly created files and previously saved files. Before accessing any file, you must lock the file and obtain the file key. A new file key is used by the locking mechanism for each lock request. The key is discarded when the lock timeout of 15 seconds expires. The lock timeout value cannot be changed.

Port Swapping

The FICON port-swapping feature is only provided for maintenance purposes.

The FICON port-swapping feature causes all configurations associated with *old-port-number* and *new port-number* to be swapped, including VSAN configurations.

Cisco MDS switches allow port swapping for nonexistent ports as follows:

- Only FICON-specific configurations (prohibit, block, and port address mapping) are swapped.
- No other system configuration is swapped.
- All other system configurations are only maintained for existing ports.
- If you swap a port in a module that has unlimited oversubscription ratios enabled with a port in a module that has limited oversubscription ratios, then you may experience a degradation in bandwidth.



Tip

If you check the **Active=Saved** check box **active equals saved** is enabled on any FICON VSAN, then the swapped configuration is automatically saved to startup. Otherwise, you must explicitly save the running configuration immediately after swapping the ports.

Once you swap ports, the switch automatically performs the following actions:

- Shuts down both the old and new ports.
- Swaps the port configuration.

If you attempt to bring the port up, you must explicitly shut down the port to resume traffic.



Note

To view the latest FICON information, you must click the Refresh button. See the [Automatically Saving the Running Configuration, on page 551](#).

FICON Tape Acceleration

The sequential nature of tape devices causes each I/O operation to the tape device over an FCIP link to incur the latency of the FCIP link. Throughput drastically decreases as the round-trip time through the FCIP link increases, leading to longer backup windows. Also, after each I/O operation, the tape device is idle until the next I/O arrives. Starting and stopping of the tape head reduces the lifespan of the tape, except when I/O operations are directed to a virtual tape.

Cisco MDS NX-OS software provides acceleration for the following FICON tape write operations:

- The link between mainframe and native tape drives (both IBM and Sun/STK)
- The back-end link between the VSM (Virtual Storage Management) and tape drive (Sun/STK)

FICON tape acceleration over FCIP provides the following advantages:

- Efficiently utilizes the tape device by decreasing idle time
- More sustained throughput as latency increases
- Similar to FCP tape acceleration, and does not conflict with it



Note

FICON tape read acceleration over FCIP is supported from Cisco MDS NX-OS Release 5.0(1). For more information refer to the [Configuring FICON Tape Read Acceleration, on page 556](#).

Figure 70: Host Directly Accessing IBM/STK (StorageTek) Library, on page 539 through Figure 73: Host Accessing Peer-to-Peer VTS (Virtual Tape Server), on page 540 show supported configurations.

Figure 70: Host Directly Accessing IBM/STK (StorageTek) Library



Figure 71: Host Accessing Standalone IBM-VTS (Virtual Tape Server)/STK-VSM (Virtual Shared Memory)



Figure 72: Host Accessing Peer-to-Peer VTS (Virtual Tape Server)





Figure 73: Host Accessing Peer-to-Peer VTS (Virtual Tape Server)



Note For information about FCIP tape acceleration, refer to the *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*.

CUP In-Band Management

The CUP protocol configures access control and provides unified storage management capabilities from a mainframe computer. Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for in-band management using the IBM S/A OS/390 I/O operations console.



Note The CUP specification is proprietary to IBM.

CUP is supported by switches and directors in the Cisco MDS 9000 Family. The CUP function allows the mainframe to manage the Cisco MDS switches.

Host communication includes control functions such as blocking and unblocking ports, as well as monitoring and error reporting functions.

Licensing Requirements for FICON

The following table shows the licensing requirements for this feature:

License	License Description
MAINFRAME_PKG	The mainframe license is required to enable FICON. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .
XRC_ACCL	The Extended Remote Copy (XRC) acceleration is required to activate FICON XRC acceleration on the Cisco MDS 9222i Switch and on the MSM-18/4 module in the Cisco MDS 9500 Series directors. For a complete explanation of the NX-OS licensing scheme and how to obtain and apply licenses, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

FICON Port Numbering Guidelines

The following guidelines apply to FICON port numbers:

- Supervisor modules do not have port number assignments.
- Port numbers do not change based on TE ports. Since TE ports appear in multiple VSANs, chassis-wide unique port numbers should be reserved for TE ports.
- Each PortChannel must be explicitly associated with a FICON port number.
- When the port number for a physical PortChannel becomes uninstalled, the relevant PortChannel configuration is applied to the physical port.
- Each FCIP tunnel must be explicitly associated with a FICON port number. If the port numbers are not assigned for PortChannels or for FCIP tunnels, then the associated ports will not come up.

See the [About Port Numbers for FCIP and PortChannel](#), on page 533.

Port Swapping Guidelines

Be sure to follow these guidelines when using the FICON port swapping feature:

- Port swapping is not supported for logical ports (PortChannels, FCIP links). Neither the *old-port-number* nor the *new-port-number* can be a logical port.
- Port swapping is not supported between physical ports that are part of a PortChannel. Neither the *old-port-number* nor the *new-port-number* can be a physical port that is part of a PortChannel.
- Before performing a port swap, the Cisco NX-OS software performs a compatibility check. If the two ports have incompatible configurations, the port swap is rejected with an appropriate reason code. For example, if a port with BB_credits as 25 is being swapped with an OSM port for which a maximum of 12 BB_credits is allowed (not a configurable parameter), the port swapping operation is rejected.
- Before performing a port swap, the Cisco NX-OS software performs a compatibility check to verify the extended BB_credits configuration.
- If ports have default values (for some incompatible parameters), then a port swap operation is allowed and the ports retain their default values.
- Port tracking information is not included in port swapping. This information must be configured separately (refer to the *Cisco MDS 9000 Family NX-OS Quality of Service Configuration Guide*).

**Note**

The 32-port module guidelines also apply for port swapping configurations (Refer to the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide*).

FICON Tape Acceleration Configuration Guidelines

FICON tape acceleration has the following configuration guidelines:

- In addition to the normal FICON configuration, FICON tape acceleration must be enabled on both ends of the FCIP interface. If only one end has FICON tape acceleration enabled, acceleration does not occur.
- FICON tape acceleration is enabled on a per VSAN basis.
- FICON tape acceleration cannot function if multiple ISLs are present in the same VSAN (PortChannels or FSPF load balanced).

- You can enable both Fibre Channel write acceleration and FICON tape acceleration on the same FCIP interface.

Enabling or disabling FICON tape acceleration disrupts traffic on the FCIP interface.

Default Settings

[Table 65: Default FICON Settings](#), on page 542 lists the default settings for FICON features.

Table 65: Default FICON Settings

Parameters	Default
FICON feature	Disabled.
Port numbers	Same as port addresses.
FC ID last byte value	0 (zero).
EBCDIC format option	US-Canada.
Switch offline state	Hosts are allowed to move the switch to an offline state.
Mainframe users	Allowed to configure FICON parameters on Cisco MDS switches.
Clock in each VSAN	Same as the switch hardware clock.
Host clock control	Allows host to set the clock on this switch.
SNMP users	Configure FICON parameters.
Port address	Not blocked.
Prohibited ports	Ports 90–253 and 255 for the Cisco MDS 9200 Series switches. Ports 250–253 and 255 for the Cisco MDS 9500 Series switches.

Configuring FICON

By default FICON is disabled in all switches in the Cisco MDS 9000 Family. You can enable FICON on a per VSAN basis by using the Device Manager.

Assigning FICON Port Numbers to Slots

**Caution**

When you assign, change, or release a port number, the port reloads.

To assign FICON port numbers to slots using Device Manager, follow these steps:

Procedure

- Step 1** Click **FICON** and then select **Port Numbers**.
You see the FICON port number.
- Step 2** Enter the chassis slot port numbers in the Reserved Port Numbers field.
- Step 3** Click **Apply**.
-

Reserving FICON Port Numbers for FCIP and PortChannel Interfaces

You must reserve port numbers for logical interfaces, such as FCIP and PortChannels, if you plan to use them. To reserve FICON port numbers for FCIP and PortChannel interfaces using Device Manager, follow these steps:

Procedure

- Step 1** Click **FICON > Port Numbers**.
You see the FICON port numbers dialog box.
- Step 2** Click the **Logical** tab to see the reserved port numbers for the slot.
- Step 3** Enter the chassis slot port numbers. These are the reserved port numbers for one chassis slot. There can be up to 64 port numbers reserved for each slot in the chassis.
- Step 4** Click **Apply**.
-

Setting Up a Basic FICON Configuration

This section steps you through the procedure to set up FICON on a specified VSAN in a Cisco MDS 9000 Family switch.



Note Press **Ctrl-C** at any prompt to skip the remaining configuration options and proceed with what is configured until that point.



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

To enable and set up FICON, follow these steps:

Procedure

Step 1 Enter the **setup ficon** command at the EXEC command mode.

Example:

```
switch# setup ficon
      --- Ficon Configuration Dialog ---
This setup utility will guide you through basic Ficon Configuration
on the system.
Press Enter if you want to skip any dialog. Use ctrl-c at anytime
to skip all remaining dialogs.
```

Step 2 Enter **yes** (the default is **yes**) to enter the basic FICON configuration setup.

Example:

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

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

Step 3 Enter the VSAN number for which FICON should be enabled.

Example:

```
Enter vsan [1-4093]:2
```

Step 4 Enter **yes** (the default is **yes**) to create a VSAN.

Example:

```
vsan 2 does not exist, create it? (yes/no) [yes]: yes
```

Step 5 Enter **yes** (the default is **yes**) to confirm your VSAN choice:

Example:

```
Enable ficon on this vsan? (yes/no) [yes]: yes
```

Note At this point, the software creates the VSAN if it does not already exist.

Step 6 Enter the domain ID number for the specified FICON VSAN.

Example:

```
Configure domain-id for this ficon vsan (1-239):2
```

Step 7 Enter **yes** (the default is **no**) to set up FICON in cascaded mode. If you enter **no**, skip to [Step 8, on page 545](#) (see the [CUP In-Band Management, on page 540](#)).

Example:

```
Would you like to configure ficon in cascaded mode: (yes/no) [no]: yes
```

a) Assign the peer WWN for the FICON: CUP.

Example:

```
Configure peer wwn (hh:hh:hh:hh:hh:hh:hh:hh): 11:00:02:01:aa:bb:cc:00
```

- b) Assign the peer domain ID for the FICON: CUP

Example:

```
Configure peer domain (1-239) :4
```

- c) Enter **yes** if you wish to configure additional peers (and repeat Steps 7.a, on page 544 and 7.b, on page 545). Enter **no**, if you do wish to configure additional peers.

Example:

```
Would you like to configure additional peers: (yes/no) [no]: no
```

- Step 8** Enter **yes** (the default is **yes**) to allow SNMP permission to modify existing port connectivity parameters (see the [Configuring SNMP Control of FICON Parameters, on page 551](#)).

Example:

```
Enable SNMP to modify port connectivity parameters? (yes/no) [yes]: yes
```

- Step 9** Enter **no** (the default is **no**) to allow the host (mainframe) to modify the port connectivity parameters, if required (see the [Allowing the Host to Change FICON Port Parameters, on page 550](#)).

Example:

```
Disable Host from modifying port connectivity parameters? (yes/no) [no]: no
```

- Step 10** Enter **yes** (the default is **yes**) to enable the **active equals saved** feature (see the [Automatically Saving the Running Configuration, on page 551](#)).

Example:

```
Enable active=saved? (yes/no) [yes]: yes
```

- Step 11** Enter **yes** (the default is **yes**) if you wish to configure additional FICON VSANs.

Example:

```
Would you like to configure additional ficon vsans (yes/no) [yes]: yes
```

- Step 12** Review and edit the configuration that you have just entered.

- Step 13** Enter **no** (the default is **no**) if you are satisfied with the configuration.

Note For documentation purposes, the following configurations shows three VSANs with different FICON settings. These settings provide a sample output for different FICON scenarios.

Example:

```
The following configuration will be applied:
fcdomain domain 2 static vsan 1
fcdomain restart disruptive vsan 1
fabric-binding database vsan 1
swwn 11:00:02:01:aa:bb:cc:00 domain 4
fabric-binding activate vsan 1
zone default-zone permit vsan 1
ficon vsan 1
```

```

no host port control
fcdomain domain 3 static vsan 2
fcdomain restart disruptive vsan 2
fabric-binding activate vsan 2 force
zone default-zone permit vsan 2
ficon vsan 2
no host port control
no active equals saved
vsan database
vsan 3
fcdomain domain 5 static vsan 3
fcdomain restart disruptive vsan 3
fabric-binding activate vsan 3 force
zone default-zone permit vsan 3
ficon vsan 3
no snmp port control
no active equals saved
Would you like to edit the configuration? (yes/no) [no]: no

```

Step 14 Enter yes (the default is **yes**) to use and save this configuration. The implemented commands are displayed. After FICON is enabled for the specified VSAN, you are returned to the EXEC mode switch prompt.

Example:

```

Use this configuration and apply it? (yes/no) [yes]: yes
`fcdomain domain 2 static vsan 1`
`fcdomain restart disruptive vsan 1`
`fabric-binding database vsan 1`
`swmn 11:00:02:01:aa:bb:cc:00 domain 4`
`fabric-binding activate vsan 1`
`zone default-zone permit vsan 1`
`ficon vsan 1`
`no host port control`
`fcdomain domain 3 static vsan 2`
`fcdomain restart disruptive vsan 2`
`fabric-binding activate vsan 2 force`
`zone default-zone permit vsan 2`
`ficon vsan 2`
`no host port control`
`no active equals saved`

```

Note If a new VSAN is created, two additional commands are displayed— **vsan database** and **vsan number**.

Example:

```

`vsan database`
`vsan 3`
`in-order-guarantee vsan 3`
`fcdomain domain 2 static vsan 3`
`fcdomain restart disruptive vsan 3`
`fabric-binding activate vsan 3 force`
`zone default-zone permit vsan 3`
`ficon vsan 3`
`no snmp port control`
Performing fast copy config...done.
switch#

```


Enabling FICON on a VSAN

To create a FICON-enabled VSAN, follow these steps:

Procedure

- Step 1** Click the Create VSAN icon (see [Figure 74: Create VSAN Icon, on page 547](#)).

Figure 74: Create VSAN Icon



You see the Create VSAN dialog box.

- Step 2** Select the switches you want to be in the VSAN.
- Step 3** Enter a VSAN ID.
- Step 4** Enter the name of the VSAN, if desired.
- Step 5** Select the type of load balancing, the interop value, and the administrative state for this VSAN.
- Step 6** Check the FICON check box.
- Note** You cannot enable interop modes on FICON-enabled VSANs.
- Step 7** Check the option, if appropriate, to enable fabric binding for the selected switches.
- Step 8** Check the All Ports Prohibited option if all ports in this VSAN are prohibited.
- Step 9** Click Create to create the VSAN.
- Step 10** Choose **Tools > Device Manager** to open Device Manager for each switch in the FICON VSAN.
- Step 11** Choose FC > VSANs.
- You see the VSAN dialog box.
- Step 12** Enter the VSAN membership information.
- Step 13** Click the VSAN you want to become a FICON VSAN and select Add from the FICON drop-down menu.
- Step 14** Click Apply to save these changes.

Manually Enabling FICON on a VSAN



Note This section describes the procedure to manually enable FICON on a VSAN. If you have already enabled FICON on the required VSAN using the automated setup (recommended), skip to the [Automatically Saving the Running Configuration, on page 551](#).

To manually enable FICON on a VSAN, follow these steps:

Procedure

- Step 1** Choose **VSAN > FICON**.
You see the FICON VSAN configuration information in the Information pane.
- Step 2** Select the switch in the VSAN on which you want to enable FICON.
- Step 3** Click enable from the Command drop-down menu.
- Step 4** Click the **Apply Changes** icon to save these changes.
-

Deleting FICON VSANs

To delete a FICON VSAN, follow these steps:

Procedure

- Step 1** Select All VSANS.
You see the VSAN table in the Information pane.
- Step 2** Click anywhere in the row of the VSAN that you want to delete.
- Step 3** Click Delete Row to delete the VSAN.
- Note** Deleting the VSAN will also delete the associated FICON configuration file, and the file cannot be recovered.
-

Suspending a FICON VSAN

To suspend a FICON VSAN, follow these steps:

Procedure

- Step 1** Click **All VSANs**.
You see all the VSANs listed in the Information pane.
- Step 2** Select the VSAN that you want to suspend.
- Step 3** Set the Admin drop-down menu for a VSAN to suspended.
- Step 4** Click the Apply Changes icon to save these changes.
-

What to do next



Note This command can be issued by the host if the host is allowed to do so (see the [Allowing the Host to Move the Switch Offline, on page 550](#)).

Configuring the code-page Option

FICON strings are coded in Extended Binary-Coded Decimal Interchange Code (EBCDIC) format. Refer to your mainframe documentation for details on the code-page options.

Cisco MDS switches support **international-5**, **france**, **brazil**, **germany**, **italy**, **japan**, **spain-latinamerica**, **uk**, and **us-canada** (default) EBCDIC format options.

To modify the code-page option using Device Manager, follow these steps:

Procedure

- Step 1** Choose **FICON > VSANs**.
You see the FICON VSAN configuration dialog box. The VSANs tab is the default tab.
- Step 2** From the CodePage drop-down menu, choose an option for the FICON VSAN you want to configure.
- Step 3** Click Apply to save the changes.

Assigning FC ID Last Byte



Note If the FICON feature is configured in cascaded mode, the Cisco MDS switches use ISLs to connect to other switches.

To assign the last byte for the FC ID, follow these steps:

Procedure

- Step 1** Choose **All VSANs > Domain Manager**.
- Step 2** Click the **Persistent FCIDs** tab.
- Step 3** Select **single** in the Mask column and then assign the entire FC ID at once. The single option allows you to enter the FC ID in the ##### format.
- Step 4** Click the **Apply Changes** icon to save these changes.

Allowing the Host to Move the Switch Offline

By default, hosts are allowed to move the switch to an offline state. To do this, the host sends a "Set offline" command (x'FD') to the CUP.

To allow the host (mainframe) to move the switch to an offline state, follow these steps:

Procedure

-
- Step 1** Choose **VSAN > FICON**.
You see a list of switches under the Control tab in the Information pane.
 - Step 2** Click the **VSANs** tab.
You see the FICON VSAN configuration information in the Information pane.
 - Step 3** Check the **Host Can Offline Sw** check box to allow the mainframe to move a switch to the offline state.
 - Step 4** Check the **Host Can Sync Time** check box to allow the mainframe to set the system time on the switch.
 - Step 5** Click the **Apply Changes** icon to save the changes.
-

Allowing the Host to Change FICON Port Parameters

By default, mainframe users are not allowed to configure FICON parameters on Cisco MDS switches—they can only query the switch.

Use the **host port control** command to permit mainframe users to configure FICON parameters.

To allow the host (mainframe) to configure FICON parameters on the Cisco MDS switch, follow these steps:

Procedure

-
- Step 1** Choose **VSAN > FICON**.
You see a list of switches under the **Control** tab in the Information pane.
 - Step 2** Click the **VSANs** tab.
You see the FICON VSAN configuration information in the Information pane.
 - Step 3** Check the **Port Control By Host** check box to allow the mainframe to control a switch.
 - Step 4** Click the **Apply Changes** icon to save the changes.
-

Allowing the Host to Control the Timestamp

By default, the clock in each VSAN is the same as the switch hardware clock. Each VSAN in a Cisco MDS 9000 Family switch represents a virtual director. The clock and time present in each virtual director can be different. To maintain separate clocks for each VSAN, the Cisco NX-OS software maintains the difference of the VSAN-specific clock and the hardware-based director clock. When a host (mainframe) sets the time, the Cisco NX-OS software updates this difference between the clocks. When a host reads the clock, it computes

the difference between the VSAN-clock and the current director hardware clock and presents a value to the mainframe.

The VSAN-clock current time is reported in the output of **show ficon vsan** *vsan-id*, **show ficon**, and **show accounting log** commands.

To configure host (mainframe) control for the VSAN time stamp, follow these steps:

Procedure

-
- Step 1** Choose **VSAN > FICON**.
You see a list of switches under the Control tab in the Information pane.
- Step 2** Click the **VSANs** tab.
You see the FICON VSAN configuration information in the Information pane.
- Step 3** Check the Host Can Sync Time checkbox to allow the mainframe to set the system time on the switch.
- Step 4** Click the Apply Changes icon to save these changes.
-

Configuring SNMP Control of FICON Parameters

By default, SNMP users can configure FICON parameters using Cisco DCNM for SAN.



- Note** If you disable SNMP in the Cisco MDS switch, you cannot configure FICON parameters using DCNM-SAN. To configure SNMP control of FICON parameters, follow these steps:
-

Procedure

-
- Step 1** Choose **VSAN > FICON**.
You see a list of switches under the Control tab in the Information pane.
- Step 2** Click the **VSANs** tab.
You see the FICON VSAN configuration information in the Information pane.
- Step 3** Check the **Port Control** By SNMP checkbox to allow SNMP users to configure FICON on the switch.
- Step 4** Click the Apply Changes icon to save these changes.
-

Automatically Saving the Running Configuration

To save the running configuration, follow these steps:

Procedure

- Step 1** Choose **VSAN > FICON**.
You see a list of switches under the Control tab in the Information pane.
- Step 2** Click the **VSANs** tab.
You see the FICON VSAN configuration information in the Information pane.
- Step 3** Check the **Active=Saved** check box to automatically save the running configuration to the startup configuration whenever there is a FICON configuration change.
- Step 4** Click the **Apply Changes** icon to save these changes.
-

Configuring FICON Ports

You can perform FICON configurations on a per-port address basis in the Cisco MDS 9000 Family switches. Even if a port is uninstalled, the port address-based configuration is accepted by the Cisco MDS switch. This configuration is applied to the port when the port becomes installed.

Binding Port Numbers to PortChannels



Caution

All port number assignments to PortChannels or FCIP interfaces are lost (cannot be retrieved) when FICON is disabled on all VSANs.

You can bind (or associate) a PortChannel with a FICON port number to bring up that interface.

Binding Port Numbers to FCIP Interfaces

You can bind (or associate) an FCIP interface with a FICON port number to bring up that interface.

Configuring Port Blocking

If you block a port, the port is retained in the operationally down state. If you unblock a port, a port initialization is attempted. When a port is blocked, data and control traffic are not allowed on that port.

Physical Fibre Channel port blocks will continue to transmit an Off-line state (OLS) primitive sequence on a blocked port.



Note

The **shutdown/no shutdown** port state is independent of the **block/no block** port state.



Note You cannot block or prohibit the CUP port (0XFE). If a port is shut down, unblocking that port does not initialize the port.

To block or unblock port addresses in a VSAN using Device Manager, follow these steps:

Procedure

- Step 1** Choose **FICON > VSANs**.
 - Step 2** Select a VSAN ID and click **Port Configuration**.
You see the FICON Port Configuration dialog box for the selected VSAN.
 - Step 3** Check the **Blocked** check box for the port that you want to block.
 - Step 4** Click **Apply** to save the changes.
-

Configuring the Default State for Port Prohibiting

By default, port prohibiting is disabled on the implemented interfaces on the switch. As of Cisco MDS SAN-OS Release 3.0(2), you can change the default port prohibiting state to enabled in VSANs that you create and then selectively disable port prohibiting on implemented ports, if desired. Also, only the FICON configuration files created after you change the default have the new default setting (see the [FICON Configuration Files, on page 537](#)).

Configuring Port Prohibiting

To prohibit port addresses in a VSAN using Device Manager, follow these steps:

Procedure

- Step 1** Choose **FICON > VSANs**.
 - Step 2** Select a VASAN ID and click Port Configuration.
You see the FICON Port Configuration dialog box.
 - Step 3** Set the port prohibit configuration for the selected FICON VSANs.
 - Step 4** Click Apply to save these changes.
-

Assigning a Port Address Name



Note To view the latest FICON information, you must click the Refresh button. See the [Automatically Saving the Running Configuration, on page 551](#).

To assign a port address name in Device Manager, follow these steps:

Procedure

-
- Step 1** Choose **FICON > VSANs**.
- Step 2** Select a VSAN ID and click Port Configuration.
You see the FICON Port Configuration dialog box.
- Step 3** Enter the Port Configuration information.
- Step 4** Click Apply to save the configuration information.
-

Specifying an RLIR Preferred Host

As of Cisco MDS SAN-OS Release 3.0(3), you can specify a preferred host to receive RLIR frames. The MDS switch sends RLIR frames to the preferred host only if it meets the following conditions:

- No host in the VSAN is registered for RLIR with the registration function set to “always receive.” If one or more hosts in the VSAN are registered as “always receive,” then RLIR sends only to these hosts and not to the configured preferred host.
- The preferred host is registered with the registration function set to “conditionally receive.”



Note If all registered hosts have the registration function set to “conditionally receive,” then the preferred host receives the RLIR frames.

You can specify only one RLIR preferred host per VSAN. By default, the switch sends RLIR frames to one of the hosts in the VSAN with the register function set to “conditionally receive” if no hosts have the register function set to “always receive.”

Clearing RLIR Information

Use the **clear rlir statistics** command to clear all existing statistics for a specified VSAN.

```
switch# clear rlir statistics vsan 1
```

Use the **clear rlir history** command to clear the RLIR history where all link incident records are logged for all interfaces.

```
switch# clear rlir history
```

Use the **clear rlir recent interface** command to clear the most recent RLIR information for a specified interface.

```
switch# clear rlir recent interface fc 1/2
```

Use the **clear rlir recent portnumber** command to clear the most recent RLIR information for a specified port number.

```
switch# clear rlir recent portnumber 16
```

Applying the Saved Configuration Files to the Running Configuration

To apply the saved configuration files to the running configuration using Device Manager, follow these steps:

Procedure

-
- Step 1** Choose **FICON > VSANs**.
 - Step 2** Click the Files tab.
You see the FICON Files dialog box.
 - Step 3** Highlight the file you want to apply and click Apply **File** to apply the configuration to the running configuration.
-

Editing FICON Configuration Files

The configuration file submode allows you to create and edit FICON configuration files. If a specified file does not exist, it is created. Up to 16 files can be saved. Each file name is restricted to eight alphanumeric characters.



Note To view the latest FICON information, you must click the Refresh button. See the [Automatically Saving the Running Configuration, on page 551](#).

To edit the contents of a specified FICON configuration file using Device Manager, follow these steps:

Procedure

-
- Step 1** Choose **FICON > VSANs**.
 - Step 2** Click the Files tab.
You see the FICON VSANs dialog box.
 - Step 3** Select a VSAN ID and then click Open to edit the FICON configuration file.
 - Step 4** Select a VSAN ID and then click Delete to delete the FICON configuration file.
 - Step 5** Click Apply to apply the changed FICON configuration file.
-

Copying FICON Configuration Files

To copy an existing FICON configuration file using Device Manager, follow these steps:

Procedure

-
- Step 1** Choose **FICON > VSANs**.
 - Step 2** Click the Files tab.

You see the FICON VSANs dialog box.

Step 3 Click Create to create a FICON configuration file.

You see the Create FICON VSANs Files dialog box.

- a) Select a VSAN ID for the FICON VSAN you want to configure.
- b) Enter the file name and the description.
- c) Click Create to create the file.

Step 4 Click Copy to copy the file to a new file.

Step 5 Click Apply to apply the FICON configuration file.

Swapping Ports

To swap ports using Device Manager, follow these steps:

Procedure

Step 1 Select two Fibre Channel ports by holding down the CTRL key and clicking them.

Step 2 Choose **FICON > Swap Selected Ports**.

Configuring FICON Tape Acceleration

To configure FICON tape acceleration over FCIP, follow these steps:

Procedure

Step 1 Expand **ISL** and then select **FCIP** in the Physical Attributes pane.

Step 2 Click the **Tunnels** tab in the Information pane.

You see a list of available switches.

Step 3 Click the **Create Row** icon to create an FCIP tunnel.

You see the Create FCIP Tunnel dialog box.

Step 4 Configure the tunnel with the options.

Step 5 Check the **TapeAccelerator** check box to enable FICON tape acceleration over this FCIP tunnel.

Step 6 Click **Create**.

Configuring FICON Tape Read Acceleration

All the configuration guidelines and restrictions applicable for FICON tape acceleration are also applicable for FICON tape read acceleration. Both FICON tape acceleration and FICON tape read acceleration can coexist.

Configuring XRC Acceleration

IBM z/OS Global Mirror eXtended Remote Copy (XRC) is supported on the MSM-18+4 modules. For XRC to function, XRC acceleration must be enabled on the FCIP tunnel interfaces on both ends. XRC acceleration is disabled by default.

XRC acceleration and FICON tape acceleration cannot be enabled on the same FCIP tunnel interface and cannot exist in the same VSAN.

Configure XRC Acceleration

To configure XRC acceleration on a FCIP tunnel interface, follow these steps:

Procedure

-
- | | |
|---------------|---------------------------------------------------------------------------------------------------------------------|
| Step 1 | Expand ISL and then select FCIP in the Physical Attributes pane. |
| Step 2 | Click the Tunnels(Advanced) tab in the Information pane.

You see a list of available FCIP interfaces. |
| Step 3 | Check the check box in the XRC Emulator column to enable XRC acceleration over the FCIP tunnel. |
| Step 4 | Click Apply . |
-

Configure XRC acceleration on an FCIP Tunnel Interface Using Device Manager

To configure XRC acceleration on an FCIP tunnel interface using Device Manager, follow these steps:

Procedure

-
- | | |
|---------------|-----------------------------------------------------------------------------------------------------------|
| Step 1 | In the Device Manager window, click IP and then select FCIP from the menu. |
| Step 2 | Click the Tunnels(Advanced) tab in the Information pane.

You see a list of FCIP interfaces. |
| Step 3 | Check the check box in the XRC Emulator column to enable XRC acceleration over the FCIP tunnel. |
| Step 4 | Click Apply . |
-

Place the CUP in a Zone

To place the CUP in a zone, follow these steps:

Procedure

-
- | | |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | In DCNM-SAN, choose Zone > Edit Full Zoneset, and then choose Edit > Edit Default Zone Attributes to set the default zone to permit for the required VSAN. |
| Step 2 | In Device Manager, choose FC > Name Server... for the required VSAN and obtain the FICON:CUP WWN. |

Note If more than one FICON:CUP WWN exists in this fabric, be sure to add all the FICON:CUP pWWNs to the required zone.

Step 3 In DCNM-SAN, choose Zone > Edit Full Zoneset and add the FICON:CUP pWWN to the zone database.

Calculating FICON Flow Load Balance

The FICON Flow Load Balance Calculator allows you to get the best load balancing configuration for your FICON flows. The calculator does not rely on any switch or flow discovery in the fabric. It is available from the DCNM-SAN Tools menu.

To use the FICON Flow Load Balance Calculator, follow these steps:

Procedure

Step 1 Choose Tools > Flow Load Balance Calculator.

You see the Flow Load Balance Calculator.

Step 2 Click **Add** to enter the source and destination(s) flows.

Step 3 Enter source and destination using 2 byte hex (by domain and area IDs). You can copy and paste these IDs, and then edit them if required.

Step 4 Enter (or select) the number of ISLs between the two switches (for example, between domain ID 0a and 0b).

Step 5 Select a row to remove it and click **Remove**.

Step 6 Select the module for which you are calculating the load balance.

Step 7 Click Calculate to show the recommended topology.

Note If you change flows or ISLs, you must click **Calculate** to see the new recommendation.

Receiving FICON Alerts

To receive an alert to indicate any changes in the FICON configuration using Device Manager, follow these steps:

Procedure

Step 1 Choose **FICON > VSANs**.

You see the FICON VSANs dialog box.

Step 2 Check the User Alert Mode check box to receive an alert when the FICON configuration changes.

Step 3 Click Apply to apply this change.

Viewing ESCON Style Ports

To view the available and prohibited ESCON style ports using Device Manager, follow these steps:

Procedure

- Step 1** Check the ESCON Style check box to see the available and prohibited ESCON style ports. In [Figure 75: ESCON Style](#), on page 559, A stands for available and P stands for prohibited. When the port address is highlighted red, it represents the E/TE port or multiple interfaces.

Figure 75: ESCON Style

Port Address	Name	Block	00	01	02	03	04	05	06	07	20	21	2D
00(fc1/1)	host1		A	A	A	A	A	A	A	A	P	A	P
01(fc1/2)		B	A	A	A	A	A	A	A	A	P	A	A
02(fc1/3)	host2		A	A	A	A	A	A	A	A	P	P	A
03(fc1/4)	host1		A	A	A	A	A	A	A	A	P	A	P
04(fc1/5)	host1		A	A	A	A	A	A	A	A	P	A	A
05(fc1/6)	host1		A	A	A	A	A	A	A	A	P	A	A
06(fc1/7)	host1		A	A	A	A	A	A	A	A	P	A	A
07(fc1/8)	storage1		P	P	P	P	P	P	P	P	P	P	P
20(fc2/1)	storage2		A	A	P	A	A	A	A	A	P	A	A
21(fc2/2)	storage3		A	A	A	P	A	A	A	A	P	A	A
2D(fc2/14)	storage4		P	A	A	A	A	A	A	A	P	A	A

11 row(s)

Buttons: Apply, Refresh, Help, Close

- Step 2** Click Apply to save the changes.

Displaying RLIR Information

To view RLIR information using Device Manager, follow these steps:

Procedure

- Step 1** Choose FICON > RLIR ERL.
You see the Show RLIR ERL dialog box.
- Step 2** Click Close to close the dialog box.

Displaying FICON Configuration Files

To open and view configuration files in DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Choose **FICON > VSAN**.
You see the FICON configuration table in the Information pane.
 - Step 2** Click the **Files** tab.
 - Step 3** Select the file you want to open.
 - Step 4** Click **Open**.
-

Displaying XRC Acceleration Statistics

To display XRC acceleration statistics, follow these steps:

Procedure

-
- Step 1** Expand **ISL** and then select **FCIP** in the Physical Attributes pane.
 - Step 2** Click the **XRC Statistics** tab in the Information pane.
You see the XRC session statistics.
-

Displaying XRC Acceleration Statistics

Procedure

-
- Step 1** In the Device Manager window, click **IP**, and then select **FCIP** from the menu.
 - Step 2** Click the **XRC Statistics** tab in the Information pane.
-

Displaying FICON Port Address Information

Procedure

-
- Step 1** Choose **FICON > VSANs**.
You see the FICON VSANs dialog box.
 - Step 2** Select a VSAN ID and click Port Configuration.
You see the FICON Port Configuration dialog box.
 - Step 3**
-

Displaying IPL File Information

To display the IPL file information using Device Manager, follow these steps:

Procedure

-
- Step 1** Select VSANs from the FICON menu.
 - Step 2** Click the Files tab.
You see the FICON VSANs dialog box.
 - Step 3** Select the file that you want to view and click Open.
-

Viewing the History Buffer

In the directory history buffer, the Key Counter column displays the 32-bit value maintained by Cisco MDS switches. This value is incremented when any port changes state in that VSAN. The key counter (a 32-bit value) is incremented when a FICON-related configuration is changed. Host programs can increment this value at the start of the channel program and then perform operations on multiple ports. The director history buffer keeps a log of which port address configuration was changed for each key-counter value.

The director history buffer provides a mechanism to determine the change in the port state from the previous time when a value was contained in the key counter.

To view the directory history buffer using Device Manager, follow these steps:

Procedure

-
- Step 1** Choose **FICON > VSANs**.
You see the FICON VSANs dialog box.
 - Step 2** Click the Director History button.
You see the history buffer dialog box.
 - Step 3** Click Close to close the dialog box.
-

Field Descriptions for FICON

This section displays the field descriptions for this feature.

FICON VSANs

Field	Description
VSAN ID	Uniquely identifies a VSAN within a fabric.
Host Can Offline SW	If true, it allows the host to put the system offline.

Field	Description
Host Can Sync Time	If true, the host can set the system time.
Port Control by Host	If true, the host is allowed to alter FICON Director connectivity parameters.
Port Control by SNMP	If true, SNMP manager is allowed to alter FICON director connectivity parameters.
CUP Name	The name of the control unit device.
CUP Enable	Indicates whether the control unit device is enabled.
Domain ID	Specifies the domain ID of the switch.
CodePage	The Code Page used in this VSAN.
Character Set	Character set for the code page used in this VSAN.
Active=Saved	If true, the active to saved mode is enabled. All changes will be saved to NVRAM.
User Alert Mode	If true, FICON management stations will prompt on changes.
Device Allegiance	If CUP is in allegiance state with a channel, it cannot accept any commands from any logical paths. A CUP goes in an allegiance state when it accepts command from a channel and forms an allegiance with it until the successful completion of the channel program, at which point the CUP goes in an unlocked mode.
VSAN Time	The system time in the VSAN. This could be set either by the host or be the default global time in the FICON Director. The default global time is the local time in the FICON Director.
VSAN State	Controls the state of the ports belonging to a VSAN in the context of the FICON functionality.
VSAN Serial Number	The serial number of the FICON director for this VSAN.

FICON VSANs Files

Field	Description
Description	Configuration file description.
CUP Name	The name of the control unit device.
Status	Locked indicates no change allowed. Unlocked indicates change allowed.
LastAccessed	The time this file was last accessed.
UserAlertMode	If true, director user alert mode is enabled.

Global

Field	Description
Default Port Prohibited	Check this option to block the default port.

FICON Port Attributes

Field	Description
TypeNumber	The type number for this FICON Director.
SerialNumber	The sequence number assigned to this FICON Director during manufacturing.
Tag	This is the identifier of the peer port. <ul style="list-style-type: none"> • If the peer port's unit type is channel, then PortId will be the CHPID (Channel Path Identifier) of the channel path that contains this peer port. • If the peer port is controlUnit, then PortId will be 0. • If the peer port is fabric, then PortId will be port address of the interface on the peer switch.
FcId	The fabric Id of the other side port (initiator /target). This will be filled only in the case of Fabric ports.
Status	valid—If this information is current. old—If this information is cached. Click Clear Old Attributes to clear the cache.
Name	The FICON port name.
Manufacturer	The name of the company that manufactured this FICON Director.
ModelNumber	The model number for this FICON Director.
PlantOfMfg	The plant code that identifies the plant of manufacture of this FICON Director.
UnitType	The peer type of the port that this port is communicating. ==Channel - host ==Control Unit - disk == Fabric - ISL
Alert	Displays one of the following: <ul style="list-style-type: none"> • bitErrThreshExceeded • lossOfSignalOrSync • nosReceived • primitiveSeqTimeOut • invalidPrimitiveSeq Click Clear to acknowledge and clear this alert.

FICON Port Configuration

Field	Description
Show Installed Ports Only	If true, only physically available ports will be listed in the table.

Field	Description
ESCON Style	ESCON Style Port Configuration display is the Port Configuration table in DM displaying the ESCON Style Ports. In the table, A represents the available ports and P represents the prohibited ports.
Port/ Prohibit	Enter the FICON address of the port and the prohibited list. (This is an alternative to the table grid.)
Name	The port name of this port.
Block	If true, this port will be isolated.
Prohibit Grid	Click on the grid to add or remove the ability of ports to communicate with each other.

FICON Port Numbers

Field	Description
Module	The number of the module in the chassis.
Reserved Port Numbers (Physical)	The reserved port numbers for the module.
NumPorts	The number of ports reserved for that module.
Module Name	The name of the module.
Reserved Port Numbers (Logical)	Chassis slot port numbers. Reserved port numbers for one chassis slot. There can be up to 64 port numbers reserved for each slot in the chassis.

FICON VSANs Director History

To view the latest FICON information, you must click the Refresh button.

Field	Description
KeyCounter	The key counter.
Ports Address Changed	The list of ports that have configuration change for a value of KeyCounter.



CHAPTER 25

Creating Dynamic VSANs

- [Creating Dynamic VSANs, on page 565](#)

Creating Dynamic VSANs

This chapter includes the following topics:

Information About DPVM

Port VSAN membership on the switch is assigned on a port-by-port basis. By default each port belongs to the default VSAN.

You can dynamically assign VSAN membership to ports by assigning VSANs based on the device WWN. This method is referred to as Dynamic Port VSAN Membership (DPVM). DPVM offers flexibility and eliminates the need to reconfigure the port VSAN membership to maintain fabric topology when a host or storage device connection is moved between two Cisco MDS 9000 family switches or two ports within a switch. It retains the configured VSAN regardless of where a device is connected or moved. To assign VSANs statically, see the *Configuring and Managing VSANs* chapter.

DPVM configurations are based on port world wide name (pWWN) and node world wide name (nWWN) assignments. A DPVM database contains mapping information for each device pWWN/nWWN assignment and the corresponding VSAN. The Cisco NX-OS software checks the database during a device FLOGI and obtains the required VSAN details.

The pWWN identifies the host or device and the nWWN identifies a node consisting of multiple devices. You can assign any one of these identifiers or any combination of these identifiers to configure DPVM mapping. If you assign a combination, then preference is given to the pWWN.

DPVM uses the Cisco Fabric Services (CFS) infrastructure to allow efficient database management and distribution. DPVM uses the application driven, coordinated distribution mode and the fabric-wide distribution scope (for information about CFS, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide* .



Note

DPVM does not cause any changes to device addressing. DPVM only pertains to the VSAN membership of the device, ensuring that the host gets the same VSAN membership on any port on the switch. For example, if a port on the switch has a hardware failure, you can move the host connection to another port on the switch and you do not need to update the VSAN membership manually.



Note DPVM is not supported on FL ports. DPVM is supported only on F ports.

About DPVM Configuration

To use the DPVM feature as designed, be sure to verify the following requirements:

- The interface through which the dynamic device connects to the Cisco MDS 9000 Family switch must be configured as an F port.
- The static port VSAN of the F port should be valid (not isolated, not suspended, and in existence).
- The dynamic VSAN configured for the device in the DPVM database should be valid (not isolated, not suspended, and in existence).



Note The DPVM feature overrides any existing static port VSAN membership configuration. If the VSAN corresponding to the dynamic port is deleted or suspended, the port is shut down.

About DPVM Databases

The DPVM database consists of a series of device mapping entries. Each entry consists of a device pWWN or nWWN assignment along with the dynamic VSAN to be assigned. You can configure a maximum of 16,000 DPVM entries in the DPVM database. This database is global to the whole switch (and fabric) and is not maintained for each VSAN.

The DPVM feature uses three databases to accept and implement configurations.

- Configuration (config) database—All configuration changes are stored in the configuration database when distribution is disabled.
- Active database—The database currently enforced by the fabric.
- Pending database—All configuration changes are stored in the DPVM pending database when distribution is enabled (see [About DPVM Database Distribution, on page 567](#)).

Changes to the DPVM config database are not reflected in the active DPVM database until you activate the DPVM config database. Changes to the DPVM pending database are not reflected in the config or active DPVM database until you commit the DPVM pending database. This database structure allows you to create multiple entries, review changes, and let the DPVM config and pending databases take effect.

About Autolearned Entries

The DPVM database can be configured to automatically learn (autolearn) about new devices within each VSAN. The autolearn feature can be enabled or disabled at any time. Learned entries are created by populating device pWWNs and VSANs in the active DPVM database. The active DPVM database should already be available to enable autolearn.

You can delete any learned entry from the active DPVM database when you enable autolearn. These entries only become permanent in the active DPVM database when you disable autolearn.



Note Autolearning is only supported for devices connected to F ports. Devices connected to FL ports are not entered into the DPVM database because DPVM is not supported on FL ports.

The following conditions apply to learned entries:

- If a device logs out while autolearn is enabled, that entry is automatically deleted from the active DPVM database.
- If the same device logs multiple times into the switch through different ports, then the VSAN corresponding to last login is remembered.
- Learned entries do not override previously configured and activated entries.
- Learning is a two-part process—Enabling autolearning followed by disabling autolearning. When the **auto-learn** option is enabled, the following applies:
 - Learning currently logged-in devices—Occurs from the time learning is enabled.
 - Learning new device logins—Occurs as and when new devices log in to the switch.

About DPVM Database Distribution

Using the CFS infrastructure, each DPVM server learns the DPVM database from each of its neighboring switches during the ISL bring-up process. If you change the database locally, the DPVM server notifies its neighboring switches, and that database is updated by all switches in the fabric.

If fabric distribution is enabled, all changes to the configuration database are stored in the DPVM pending database. These changes include the following tasks:

- Adding, deleting, or modifying database entries.
- Activating, deactivating, or deleting the configuration database.
- Enabling or disabling autolearning.

These changes are distributed to all switches in a fabric when you commit the changes. You can also discard (abort) the changes at this point.



Tip You can view the contents of the DPVM pending database by using the **show dpvm pending** command.



Tip See the [Viewing the Pending Database, on page 574](#) to view the contents of the of the pending database.

If the DPVM database is available on all switches in the fabric, devices can be moved anywhere and offer the greatest flexibility. To enable database distribution to the neighboring switches, the database should be consistently administered and distributed across all switches in the fabric. The Cisco NX-OS software uses the Cisco Fabric Services (CFS) infrastructure to achieve this requirement (refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*).

About Locking the Fabric

The first action that modifies the existing configuration creates the DPVM pending database and locks the feature in the fabric. Once you lock the fabric, the following conditions apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the DPVM pending database. Modifications from this point on are made to the DPVM pending database. The DPVM pending database remains in effect until you commit the modifications to the DPVM pending database or discard (abort) the changes to the DPVM pending database.

About Copying DPVM Databases

The following circumstances may require the active DPVM database to be copied to the DPVM config database:

- If the learned entries are only added to the active DPVM database.
- If the DPVM config database or entries in the DPVM config database are accidentally deleted.


Note

If you copy the DPVM database and fabric distribution is enabled, you must commit the changes.

Licensing Requirements for VSANs

The following table shows the licensing requirements for this feature:

License	License Description
ENTERPRISE_PKG	The enterprise license is required to enable VSAN. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .

Guidelines and Limitations

This section explains the database guidelines for this feature.

A database merge refers to a union of the configuration database and static (unlearned) entries in the active DPVM database. For information about CFS merge support, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide* for detailed concepts.

When merging the DPVM database between two fabrics, follow these guidelines:

- Verify that the activation status and the autolearn status is the same in both fabrics.
- Verify that the combined number of device entries in each database does not exceed 16 K.


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.

Default Settings

[Table 66: Default DPVM Parameters](#), on page 568 lists the default settings for DPVM parameters.

Table 66: Default DPVM Parameters

Parameters	Default
DPVM	Disabled.
DPVM distribution	Enabled.
Autolearning	Disabled.

Creating DPVM


This section includes the following topics:

Configuring DPVM with the DPVM Wizard

To use the DPVM Setup Wizard to set up dynamic port VSAN membership, follow these steps:

Procedure

- Step 1** Click the DPVM Setup Wizard icon in the DCNM-SAN toolbar.



You see the Select Master Switch page.

- Step 2** Click the switch you want to be the master switch. This switch controls the distribution of the DPVM database to other switches in the fabric.
- Step 3** Click Next.
- You see the AutoLearn Current End Devices page.
- Step 4** (Optional) Click the Create Configuration From Currently Logged In End Devices check box if you want to turn on autolearning.
- Step 5** Click Next.
- You see the Edit and Activate Configuration page.
- Step 6** Verify the current or autolearned configuration. Optionally, click Insert to add more entries into the DPVM config database.
- Step 7** Click Finish to update the DPVM config database, distribute the changes using CFS, and activate the database, or click Cancel to exit the DPVM Setup Wizard without saving changes.
- Step 8** Select the switch you want to be the master switch. This switch controls the distribution of the DPVM database to other switches in the fabric.
- Step 9** Click Next.
- You see the AutoLearn Current End Devices page.
- Step 10** (Optional) Check the Create Configuration From Currently Logged In End Devices check box if you want to enable autolearning.
- Step 11** Click Next.
- You see the Edit and Activate Configuration page.
- Step 12** Verify the current or autolearned configuration. Optionally, click Insert to add more entries into the DPVM config database.

- Step 13** Click Finish to update the DPVM config database, distribute the changes using CFS, and activate the database, or click Cancel to exit the DPVM Setup Wizard without saving changes.

Configuring DPVM Config and Pending Databases

To create and populate the config and pending databases, follow these steps:

Procedure

- Step 1** Expand Fabricxx> All VSANs, and then select DPVM in the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the **CFS** tab and select a master switch by checking a check box in the Master column.
- Note** You must click the CFS tab in order to activate the other tabs.
- Step 3** Click the Config Database tab and then click the Create Row to insert a new entry.
You see the Create Config Database dialog box.
- Step 4** Choose an available WWN and VSAN combination or fill in the pWWN and Login VSAN fields.
- Step 5** Click Create to save these changes in the config or pending database or click Close to discard any unsaved changes.
- Step 6** Click the CFS tab and select the Config Action drop-down menu for the master database.
- Step 7** Select commit from the drop-down menu to distribute these changes or abort to discard the changes.

Activating DPVM Config Databases

To activate the DPVM config database, follow these steps:

Procedure

- Step 1** Expand Fabricxx> All VSANs, and then select DPVM from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the **Action** tab and set the Action drop-down menu to activate or forceActivate to activate the DPVM config database.
- Step 3** Click the CFS tab and select the Config Action drop-down menu for the master database.
- Step 4** Select commit from the drop-down menu to distribute these changes or abort to discard the changes.
- Note** To disable DPVM, you must explicitly deactivate the currently active DPVM database.

Enabling Autolearning

To enable autolearning, follow these steps:

Procedure

- Step 1** Expand Fabricxx> All VSANs, and then select DPVM from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the Actions tab and check the Auto Learn Enable check box to enable autolearning.
- Step 3** Click the CFS tab and select commit to distribute these changes or abort to discard the changes.
-

Clearing a Single Autolearned Entry

To clear a single autolearn entry, follow these steps:

Procedure

- Step 1** Expand Fabricxx> All VSANs, and then select DPVM from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the Actions tab and select clearOnWWN from the Auto Learn Clear drop-down menu.
- Step 3** Check the **clear WWN** check box next to the WWN of the autolearned entry that you want to clear.
- Step 4** Click CFS and select commit to distribute these changes or abort to discard the changes.
-

Clearing All Autolearned Entries

To clear all autolearn entries, follow these steps:

Procedure

- Step 1** Expand Fabricxx> All VSANs, and then select DPVM from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the Actions tab.
You see the DPVM Actions menu.
- Step 3** Select **clear** from the Auto Learn Clear drop-down menu.
- Step 4** Click the CFS tab and select commit to distribute these changes or abort to discard the changes.
-

Disabling DPVM Database Distribution

To disable DPVM database distribution to the neighboring switches, follow these steps:

Procedure

- Step 1** Expand Fabricxx> All VSANs, and then select DPVM from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the CFS tab and select disable from the Admin drop-down menu.
- Step 3** Click Apply Changes to save this change or click Undo Changes to discard the change.
-

Locking the Fabric

To lock the fabric and apply changes to the DPVM pending database, follow these steps:

Procedure

- Step 1** Expand Fabricxx> All VSANs, and then select DPVM from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the Config Database tab and Create Row.
You see the Create Config Database dialog box.
- Step 3** Choose an available pWWN and login VSAN.
- Step 4** Click Create to save this change to the pending database or click Close to discard any unsaved change.
-

Committing Changes

If you commit the changes made to the configuration, the configuration in the DPVM pending database are distributed to other switches. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

To commit the DPVM pending database, follow these steps:

Procedure

- Step 1** Expand Fabricxx> All VSANs, and then select DPVM from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the CFS tab and select commit from the Config Action drop-down menu.
- Step 3** Click Apply Changes to save this change or click Undo Changes to discard the change.
-

Discarding Changes

If you discard (abort) the changes made to the DPVM pending database, the configurations remain unaffected and the lock is released.

To discard the DPVM pending database, follow these steps:

Procedure

-
- Step 1** Expand Fabricxx> All VSANs, and then select DPVM from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
 - Step 2** Click the CFS tab and select abort from the Config Action drop-down menu.
 - Step 3** Click Apply Changes to save this change or click Undo Changes to discard the change.
-

Clearing a Locked Session

If you have performed a DPVM task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the DPVM pending database are discarded and the fabric lock is released.



Note The DPVM pending database is only available in the volatile directory and is subject to being discarded if the switch is restarted.

To use administrative privileges and release a locked DPVM session using DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Expand Fabricxx> All VSANs, and then select DPVM from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
 - Step 2** Click the CFS tab and select clear from the Config Action drop-down menu.
 - Step 3** Click Apply Changes to save this change or click Undo Changes to discard the change.
-

Copying DPVM Databases

To copy the currently active DPVM database to the DPVM config database, follow these steps:

Procedure

-
- Step 1** Expand Fabricxx> All VSANs, and then select DPVM in the Logical Attributes pane.
You see the DPVM configuration in the Information pane.

- Step 2** Click the Actions tab and check the CopyActive to Config check box.
- Step 3** Click the CFS tab and select commit from the Config Action drop-down menu.
-

Comparing Database Differences

To compare the currently active database entries to the DPVM config database, follow these steps:

Procedure

- Step 1** Expand Fabricxx> All VSANs, and then select DPVM from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the Active Database tab.
You see the DPVM active database in the Information pane.
- Step 3** Select Config from the Compare With drop-down menu.
You see the comparison dialog box.
- Step 4** Select Close to close the comparison dialog box.
-

Viewing the Pending Database

To view the pending database, follow these steps:

Procedure

- Step 1** Expand Fabricxx> All VSANs, and then select DPVM from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the CFS tab and set the Config View drop-down menu to pending.
- Step 3** Click Apply Changes.
- Step 4** Click the Config Database tab.
You see the pending database entries.
-

Field Descriptions for DPVM

This section describes the field descriptions for this feature.

DPVM Actions

Field	Description
Action	Helps in activating the set of bindings.
Result	Indicates the outcome of the activation.
Status	Indicates the state of activation. If true, then activation has been attempted as the most recent operation. If false, then an activation has not been attempted as the most recent operation.
CopyActive to Config	When set to copy(1), results in the active (enforced) binding database to be copied on to the configuration binding database. The learned entries are also copied.
Auto Learn Enable	Helps to learn the configuration of devices logged into the local device on all its ports and the VSANs to which they are associated.
Auto Learn Clear	Assists in clearing the autolearned entries.
Clear WWN	Represents the Port WWN (pWWN) to be used for clearing its corresponding autolearned entry.

DPVM Config Database

Field	Description
Type	Specifies the type of the corresponding instance of cdpvmLoginDev object.
WWN or Name	Represents the logging-in device.
VSAN Id	Represents the VSAN to be associated to the port on the local device on which the device represented by cdpvmLoginDev logs in.
Switch Interface	Represents the device alias.

DPVM Active Database

Field	Description
Type	Specifies the type of the corresponding instance of cdpvmEnfLoginDev.
WWN or Name	Represents the logging in device address.
VSAN Id	Represents the VSAN of the port on the local device through which the device represented by cdpvmEnfLoginDev logs in.
Interface	Represents the device alias.
IsLearnt	Indicates whether this is a learned entry or not. If true, then it is a learned entry. If false, then it is not.

Additional References

For additional information related to implementing VSANs, see the following section:

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	Cisco MDS 9000 Family Command Reference

Standards


Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	–

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	–

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none"> CISCO-DYNAMIC-PORT-VSAN-MIB 	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.htm



CHAPTER 26

Distributing Device Alias Services

- [Distributing Device Alias Services, on page 577](#)

Distributing Device Alias Services

All switches in the Cisco MDS 9000 Family support Distributed Device Alias Services (device alias) on a per-VSAN basis and on a fabric-wide basis. Device alias distribution allows you to move host bus adapters (HBAs) between VSANs without manually reentering alias names.

This chapter includes the following topics:

Information About Device Aliases

When the port WWN (pWWN) of a device must be specified to configure different features (zoning, QoS, port security) in a Cisco MDS 9000 Family switch, you must assign the correct device name each time you configure these features. An incorrect device name can cause unexpected results. You can avoid this problem if you define a user-friendly name for a port WWN and use this name in all of the configuration commands as required. These user-friendly names are referred to as *device aliases* in this chapter.

About Device Alias Modes

Device alias supports two modes: basic and enhanced mode.

- When device alias runs in the basic mode, all applications function like the applications on the Cisco SAN-OS Release 3.0 switches. When you configure the basic mode using device aliases, the application immediately expands to pWWNs. This operation continues until the mode is changed to enhanced.
- When device alias runs in the enhanced mode, all applications accept the device-alias configuration in the native format. The applications store the device alias name in the configuration and distribute it in the device alias format instead of expanding to pWWN. The applications track the device alias database changes and take actions to enforce it.

A native device-alias configuration is not accepted in the interop mode VSAN. IVR zoneset activation will fail in interop mode VSANs if the corresponding twilight zones being injected are native device alias members.

Changing Mode Settings

When the device alias mode is changed from basic to enhanced mode, the applications are informed about the change. The applications start accepting the device alias-based configuration in the native format.



Note Because the device alias was previously running in the basic mode, the applications do not have any prior native device alias configuration.

The applications check for an existing device alias configuration in the native format. If the device alias is in the native format, the applications reject the request and device alias mode cannot be changed to basic.

All native device alias configurations (both on local and remote switches) must be explicitly removed, or all device alias members must be replaced with the corresponding pWWN before changing the mode back to basic.

Device Alias Mode Distribution

If the device alias distribution is turned on, it is distributed to the other switches in the network whenever there is a change in the mode. You cannot change the mode from basic to enhanced unless all the switches are upgraded to Cisco SAN-OS Release 3.1. The device alias enhancements will not apply unless the entire fabric is upgraded to Cisco SAN-OS Release 3.1.



Note When all the switches are upgraded to Cisco SAN-OS Release 3.1, you cannot automatically convert to enhanced mode. You do not need to change to enhanced mode, you can continue working in the basic mode.

Merging Device Aliases

If two fabrics are running different device alias modes and are joined together, the device alias merge will fail. There is no automatic conversion of one mode to the other during the merge process. You will need to resolve the issue.



Note Release 3.0 switches run in basic mode.

At the application level, a merger takes place between the applications and the fabric. For example, zone merge occurs when the E port is up and the IVR/PSM/DPVM merge occurs due to CFS. This merge is completely independent of the device alias merge.

If the application running on an enhanced fabric has a native device alias configuration, the application must fail the merge. The application has to fail the merge even though the other fabric can support the native device alias-based configuration, but is running in the basic mode. You will need to resolve the issue. Once the device alias merge issue is resolved, each application must be fixed accordingly.

Resolving Merge and Device Alias Mode Mismatch

If two fabrics are running in different modes and the device alias merge fails between the fabrics, the conflict can be resolved by selecting one mode or the other. If you choose the enhanced mode, ensure that all the switches are running at least the Cisco SAN-OS Release 3.1. Otherwise, the enhanced mode cannot be turned on. If you choose the basic mode, the applications running on the enhanced fabric have to comply with the device alias merge.

The device alias merge fails because of mode mismatch, but the application merge succeeds if it does not have any native device alias configurations.

If the native device alias configuration is attempted on an application from a Release 3.1 switch, the commit must be rejected because of device alias mode mismatch on some of the applications.

**Note**

The applications should not accept any native device alias configuration over SNMP if the device alias is running in the basic mode on that particular switch.

**Note**

Confcheck is added when the enhanced mode is turned on and removed when it is turned off. Applications have to add confcheck if they have a device alias configuration in the native format. They have to remove confcheck once the configuration is removed.

Device Alias Features

Device aliases have the following features:

- The device alias information is independent of your VSAN configuration.
- The device alias configuration and distribution is independent of the zone server and the zone server database.
- You can import legacy zone alias configurations without losing data.
- The device alias application uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management and distribution. Device aliases use the coordinated distribution mode and the fabric-wide distribution scope (refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*).
- When you configure zones, IVR zones, or QoS features using device aliases, and if you display these configurations, the device aliases are automatically displayed along with their respective pWWNs.

Device Alias Requirements

Device aliases have the following requirements:

- You can only assign device aliases to pWWNs.
- The mapping between the pWWN and the device alias to which it is mapped must have a one-to-one relationship. A pWWN can be mapped to only one device alias and vice versa.
- A device alias name is restricted to 64 alphanumeric characters and may include one or more of the following characters:
 - a to z and A to Z
 - 1 to 9
 - - (hyphen) and _ (underscore)
 - \$ (dollar sign) and ^ (up caret)

Zone Aliases Versus Device Aliases

Table 67: [Comparison Between Zone Aliases and Device Aliases](#), on page 580 compares the configuration differences between zone-based alias configuration and device alias configuration.

Table 67: Comparison Between Zone Aliases and Device Aliases

Zone-Based Aliases	Device Aliases
Aliases are limited to the specified VSAN.	You can define device aliases without specifying the VSAN number. You can also use the same definition in one or more VSANs without any restrictions.
Zone aliases are part of the zoning configuration. The alias mapping cannot be used to configure other features.	Device aliases can be used with any feature that uses the pWWN.
You can use any zone member type to specify the end devices.	Only pWWNs are supported along with new device aliases such as IP addresses.
Configuration is contained within the Zone Server database and is not available to other features.	Device aliases are not restricted to zoning. Device alias configuration is available to the FCNS, zone, fcping, traceroute, and IVR applications.

Device Alias Databases

The device alias feature uses two databases to accept and implement device alias configurations:

- Effective database—The database currently used by the fabric.
- Pending database—Your subsequent device alias configuration changes are stored in the pending database.

If you modify the device alias configuration, you need to commit or discard the changes as the fabric remains locked during this period.

About Device Alias Distribution

By default, device alias distribution is enabled. The device alias feature uses the coordinated distribution mechanism to distribute the modifications to all switches in a fabric.

If you have not committed the changes and you disable distribution, then a commit task will fail.

About Creating a Device Alias

When you perform the first device alias task (regardless of which device alias task), the fabric is automatically locked for the device alias feature. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the effective database is obtained and used as the pending database. Modifications from this point on are made to the pending database. The pending database remains in effect until you commit the modifications to the pending database or discard (**abort**) the changes to the pending database.

Fabric Lock Override

If you have performed a device alias task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip

The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

About Legacy Zone Alias Configuration Conversion

You can import legacy zone alias configurations to use this feature without losing data if they follow the following restrictions:

- Each zone alias has only one member.
- The member type is pWWN.
- The name and definition of the zone alias should not be the same as any existing device alias name.

If any name conflict exists, the zone aliases are not imported.



Tip Ensure 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. At this time if you do not want to distribute the configuration to other switches in the fabric, you can perform the **abort** operation and the merge changes are completely discarded.

Guidelines and Limitations

This section explains the database guidelines for this feature.

For information about CFS merge support, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide* for detailed concepts.

When merging two device alias databases, follow these guidelines:

- Verify that two device aliases with different names are not mapped to the same pWWN.
- Verify that two different pWWNs are not mapped to the same device aliases
- Verify that the combined number of the device aliases in both databases does not exceed 8191 (8K). For example, if database N has 6000 device aliases and database M has 2192 device aliases, this merge operation will fail.

Default Settings

[Table 68: Default Device Alias Parameters , on page 581](#) lists the default settings for device alias parameters.

Table 68: Default Device Alias Parameters

Parameters	Default
Database in use	Effective database.
Database to accept changes	Pending database.
Device alias fabric lock state	Locked with the first device alias task.

Configuring Device Aliases

This section includes the following topics:

Creating Device Aliases

To lock the fabric and create a device alias in the pending database, follow these steps:

Procedure

- Step 1** Expand End Devices, and then select Device Alias in the Physical Attributes pane.
You see the device alias configuration in the Information pane.
 - Step 2** Click the Configuration tab and click the Create Row icon.
You see the Device Alias Creation dialog box.
 - Step 3** Select a switch from the drop-down menu.
 - Step 4** Complete the Alias name and pWWN fields.
 - Step 5** Click Create to create this alias or click Close to discard any unsaved changes.
-

Distributing the Device Alias Database

To enable the device alias distribution, follow these steps:

Procedure

- Step 1** Expand End Devices and then select Device Alias in the Physical Attributes pane.
You see the device alias configuration in the Information pane.
The CFS tab is the default tab.
 - Step 2** Select enable from the Global drop-down menus to enabled switch aliases.
 - Step 3** Select commit from the Config Action drop-down menu for the newly enabled switches.
 - Step 4** Click Apply Changes to commit and distribute these changes or click Undo Changes to discard any unsaved changes.
-

Committing Changes

If you commit the changes made to the pending database, the following events occur:

1. The pending database contents overwrites the effective database contents.
2. The pending database is emptied of its contents.
3. The fabric lock is released for this feature.

To commit the changes to the device alias database, follow these steps:

Procedure

- Step 1** Expand End Devices, and then select Device Alias in the Physical Attributes pane.
You see the device alias configuration in the Information pane. The CFS tab is the default tab.
- Step 2** Select enable from the Global drop-down menus to enable switch aliases.
- Step 3** Select commit from the Config Action drop-down menu for the newly enabled switches.
- Step 4** Click Apply Changes to commit and distribute these changes or click Undo Changes to discard any unsaved changes.
-

Discarding Changes

If you discard 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, follow these steps:

Procedure

- Step 1** Expand End Devices and then select Device Alias in the Physical Attributes pane.
You see the device alias configuration in the Information pane. The CFS tab is the default tab.
- Step 2** Select abort from the Config Action drop-down menu.
- Step 3** Click Apply Changes to discard the session.
-

Using Device Aliases or FC Aliases

You can change whether DCNM-SAN uses FC aliases or global device aliases from DCNM-SAN Client without restarting Cisco DCNM for SAN.

To change whether DCNM-SAN uses FC aliases or global device aliases, follow these steps:

Procedure

- Step 1** Click Server > Admin.
You see the Admin dialog box.
- Step 2** For each fabric that you are monitoring with Cisco DCNM for SAN, check the Device Alias check box to use global device aliases, or uncheck to use FC aliases.
- Step 3** Click Apply to save these changes or click Close to exit the dialog box without saving any changes.
-

Populating Device Alias to Interface Description

When an end device is not logged into the switch, the Device Alias is blank. To find out what device is supposed to connect to an FC port when the device is logged out, you can populate the interface description with the device alias when the devices are logged in.

To populate the interface description with the device alias, follow these steps:

Procedure

- Step 1** From the Physical Attributes pane, expand End Devices.
 - Step 2** From the right pane, click the **General** tab.
 - Step 3** Select the rows of FC interfaces.
 - Step 4** Click the **Alias->Description** button.
 - Step 5** Click the **commit** button.
-

Rename Device Alias

There are two options for renaming the device alias form SAN client.

To rename the device alias, follow these steps:

Procedure

- Step 1** From the **Logical Domains** pane, select **Data Center > SAN > Fabric name**.
- Step 2** In the information pane on the right-hand side, click **Host Ports** or **Storage Ports**.
- Step 3** Double click the **Device Alias** column and enter the new name.
- Step 4** Click the **Apply Changes** icon.

Alternatively, you can perform the following steps to rename the device alias form SAN client.

- a. From the **Physical Attributes** pane, expand **End Devices** and select **Hosts** or **Storage**.
 - b. Double click the **Device Alias** column and enter the new name.
 - c. Click the **Apply Changes** icon.
-

Field Descriptions for Device Aliases

This section displays the field descriptions for this feature.

Device Alias Configuration

Field	Description
Device Alias	The device alias of this entry. A device can have only one alias configured.
WWN	The Fibre Channel device which is given a device alias.

Device Alias Mode

Field	Description
ConfigMode	Specifies the mode in which the device aliases can be configured. When it is set to basic, the device aliases operate in basic mode of operation. When basic mode is turned on, all MIBs which are using device aliases should internally convert them to their equivalent pWWNs and use the pWWNs. The mechanism to be followed for this conversion is implementation specific. When it is set to enhanced, the Device aliases operate in enhanced mode of operation. When enhanced mode is turned on, all MIBs which are using device aliases should use them as is without any conversion. Since the device aliases are used directly without any conversion, this is the native mode of operation of device aliases.

Device Alias Discrepancies

Field	Description
Discrepancy	Represents the checksum computed over the database represented by cfdaConfigTable and the cfdaConfigMode object. This object is used by a network manager to check if the above mentioned objects have changed on the local device. The method used to compute the checksum is implementation specific.



CHAPTER 27

Configuring Advanced Fabric Features

- [Configuring Advanced Fabric Features, on page 587](#)

Configuring Advanced Fabric Features

This chapter describes the advanced features provided in switches in the Cisco MDS 9000 Series.

Information About Common Information Model

Common Information Model (CIM) is an object-oriented information model that extends the existing standards for describing management information in a network/enterprise environment.



Note

CIM is not supported in Cisco MDS NX-OS Release 5.2(1), but is supported in Cisco DCNM Release 5.2(1).

CIM messages are independent of platform and implementation because they are encoded in N Extensible Markup Language (XML). CIM consists of a specification and a schema. The specification defines the syntax and rules for describing management data and integrating with other management models. The schema provides the actual model descriptions for systems, applications, networks, and devices.

For more information about CIM, refer to the specification available through the Distributed Management Task Force (DMTF) website at the following URL: <http://www.dmtf.org/>

For further information about Cisco MDS 9000 Family support for CIM servers, refer to the *Cisco MDS 9000 Family CIM Programming Reference Guide*.

A CIM client is required to access the CIM server. The client can be any client that supports CIM.

SSL Certificate Requirements and Format

To limit access to the CIM server to authorized clients, you can enable the HTTPS transport protocol between the CIM server and client. On the switch side, you must install a Secure Socket Library (SSL) certificate generated on the client and enable the HTTPS server. Certificates may be generated using third-party tools, such as openssl (available for UNIX, Mac, and Windows), and may be certified by a CA or self-signed.

The SSL certificate that you install on the switch must meet the following requirements:

- The certificate file contains the certificate and the private key.

- The private key must be RSA type.
- The certificate file should be in Private Electronic Mail (PEM) style format and have .pem as the extension.

```
-----BEGIN CERTIFICATE-----
(certificate goes here)
-----END CERTIFICATE-----
-----BEGIN RSA PRIVATE KEY-----
(private key goes here)
-----END RSA PRIVATE KEY-----
```

Only one certificate file can be installed at a time.

Fibre Channel Time-Out Values

You can modify Fibre Channel protocol related timer values for the switch by configuring the following time-out values (TOVs):

- Distributed services TOV (D_S_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 5,000 milliseconds.
- Error detect TOV (E_D_TOV)—The valid range is from 1,000 to 10,000 milliseconds. The default is 2,000 milliseconds. This value is matched with the other end during port initialization.
- Resource allocation TOV (R_A_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 10,000 milliseconds. This value is matched with the other end during port initialization.



Note

The fabric stability TOV (F_S_TOV) constant cannot be configured.

About fctimer Distribution

You can enable per-VSAN fctimer fabric distribution for all Cisco MDS switches in the fabric. When you perform fctimer configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you issue the first configuration command after you enabled distribution in a switch. The fctimer application uses the effective and pending database model to store or commit the commands based on your configuration.

Refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide* for more information on the CFS application.

Fabric Lock Override

If you have performed a fctimer fabric task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip

The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

World Wide Names

The world wide name (WWN) in the switch is equivalent to the Ethernet MAC address. As with the MAC address, you must uniquely associate the WWN to a single device. The principal switch selection and the allocation of domain IDs rely on the WWN. The WWN manager, a process-level manager residing on the switch's supervisor module, assigns WWNs to each switch.

Cisco MDS 9000 Family switches support three network address authority (NAA) address formats.

Table 69: 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 NX-OS software release.

Both ELPs and EFPs use the VSAN WWN by default during link initialization. However, the ELP usage changes based on the peer switch's usage:

- If the peer switch ELP uses the switch WWN, then the local switch also uses the switch WWN.
- If the peer switch ELP uses the VSAN WWN, then the local switch also uses the VSAN WWN.



Note

As of Cisco SAN-OS Release 2.0(2b), the ELP is enhanced to be compliant with FC-SW-3.

FC ID Allocation for HBAs

Fibre Channel standards require a unique FC ID to be allocated to an N port attached to a Fx port in any switch. To conserve the number of FC IDs used, Cisco MDS 9000 Family switches use a special allocation scheme.

Some HBAs do not discover targets that have FC IDs with the same domain and area. Prior to Cisco SAN-OS Release 2.0(1b), the Cisco SAN-OS software maintained a list of tested company IDs that do not exhibit this behavior. These HBAs were allocated with single FC IDs, and for others a full area was allocated.

The FC ID allocation scheme available in Release 1.3 and earlier, allocates a full area to these HBAs. This allocation isolates them to that area and are listed with their pWWN during a fabric login. The allocated FC IDs are cached persistently and are still available in Cisco SAN-OS Release 2.0(1b).

To allow further scalability for switches with numerous ports, the Cisco NX-OS software maintains a list of HBAs exhibiting this behavior. Each HBA is identified by its company ID (also known as Organizational Unique Identifier, or OUI) used in the pWWN during a fabric login. 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 kind (whole area or single) of FC ID allocated, the FC ID entries remain persistent.

Default Company ID List

All switches in the Cisco MDS 9000 Family that ship with Cisco SAN-OS Release 2.0(1b) or later, or NX-OS 4.1(1) 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.

Switch Interoperability

Interoperability enables the products of multiple vendors to interact with each other. Fibre Channel standards guide vendors towards common external Fibre Channel interfaces.

If all vendors followed the standards in the same manner, then interconnecting different products would become a trivial exercise. However, not all vendors follow the standards in the same way, thus resulting in interoperability modes. This section briefly explains the basic concepts of these modes.

Each vendor has a regular mode and an equivalent interoperability mode, which specifically turns off advanced or proprietary features and provides the product with a more amiable standards-compliant implementation.



Note

For more information on configuring interoperability for the Cisco MDS 9000 Family switches, refer to the *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*.

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, refer to the *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*.

[Table 70: Changes in Switch Behavior When Interoperability Is Enabled](#), on page 591 lists the changes in switch behavior when you enable interoperability mode. These changes are specific to switches in the Cisco MDS 9000 Family while in interop mode.

Table 70: Changes in Switch Behavior When Interoperability Is Enabled

Switch Feature	Changes if Interoperability Is Enabled
Domain IDs	Some vendors cannot use the full range of 239 domains within a fabric. Domain IDs are restricted to the range 97-127. This is to accommodate McData's nominal restriction to this same range. They can either be set up statically (the Cisco MDS switch accept only one domain ID, if it does not get that domain ID it isolates itself from the fabric) or preferred. (If it does not get its requested domain ID, it accepts any assigned domain ID.)
Timers	All Fibre Channel timers must be the same on all switches as these values are exchanged by E ports when establishing an ISL. The timers are F_S_TOV, D_S_TOV, E_D_TOV, and R_A_TOV.
F_S_TOV	Verify that the Fabric Stability Time Out Value timers match exactly.
D_S_TOV	Verify that the Distributed Services Time Out Value timers match exactly.
E_D_TOV	Verify that the Error Detect Time Out Value timers match exactly.
R_A_TOV	Verify that the Resource Allocation Time Out Value timers match exactly.
Trunking	Trunking is not supported between two different vendor's switches. This feature may be disabled on a per port or per switch basis.
Default zone	The default zone behavior of permit (all nodes can see all other nodes) or deny (all nodes are isolated when not explicitly placed in a zone) may change.

Switch Feature	Changes if Interoperability Is Enabled
Zoning attributes	<p>Zones may be limited to the pWWN and other proprietary zoning methods (physical port number) may be eliminated.</p> <p>Note Brocade uses the cfgsave command to save fabric-wide zoning configuration. This command does not have any effect on Cisco MDS 9000 Family switches if they are part of the same fabric. You must explicitly save the configuration on each switch in the Cisco MDS 9000 Family.</p>
Zone propagation	<p>Some vendors do not pass the full zone configuration to other switches, only the active zone set gets passed.</p> <p>Verify that the active zone set or zone configuration has correctly propagated to the other switches in the fabric.</p>
VSAN	<p>Interop mode only affects the specified VSAN.</p> <p>Note Interop modes cannot be enabled on FICON-enabled VSANs.</p>
TE ports and PortChannels	TE ports and PortChannels cannot be used to connect Cisco MDS to non-Cisco MDS switches. Only E ports can be used to connect to non-Cisco MDS switches. TE ports and PortChannels can still be used to connect an Cisco MDS to other Cisco MDS switches even when in interop mode.
FSPF	The routing of frames within the fabric is not changed by the introduction of interop mode. The switch continues to use src-id, dst-id, and ox-id to load balance across multiple ISL links.
Domain reconfiguration disruptive	This is a switch-wide impacting event. Brocade and McData require the entire switch to be placed in offline mode and/or rebooted when changing domain IDs.
Domain reconfiguration nondisruptive	This event is limited to the affected VSAN. Only Cisco MDS 9000 Family switches have this capability—only the domain manager process for the affected VSAN is restarted and not the entire switch.
Name server	Verify that all vendors have the correct values in their respective name server database.
IVR	IVR-enabled VSANs can be configured in no interop (default) mode or in any of the interop modes.

Guidelines and Limitations

This section explains the database merge guidelines for this feature.

When merging two fabrics, follow these guidelines:

- Be aware of the following merge conditions:
 - The merge protocol is not implemented for distribution of the fctimer values—you must manually merge the fctimer values when a fabric is merged. The per-VSAN fctimer configuration is distributed in the physical fabric.

- The fctimer configuration is only applied to those switches containing the VSAN with a modified fctimer value.
- The global fctimer values are not distributed.
- Do not configure global timer values when distribution is enabled.



Note The number of pending fctimer configuration operations cannot be more than 15. At that point, you must commit or abort the pending configurations before performing any more operations.

For information about CFS merge support, refer the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide* .

Default Settings

[Table 71: Default Settings for Advanced Features , on page 593](#) lists the default settings for the features included in this chapter.

Table 71: Default Settings for Advanced Features

Parameters	Default
CIM server	Disabled
CIM server security protocol	HTTP
D_S_TOV	5,000 milliseconds.
E_D_TOV	2,000 milliseconds.
R_A_TOV	10,000 milliseconds.
Timeout period to invoke fctrace	5 seconds.
Number of frame sent by the fcping feature	5 frames.
Remote capture connection protocol	TCP.
Remote capture connection mode	Passive.
Local capture frame limit s	10 frames.
FC ID allocation mode	Auto mode.
Loop monitoring	Disabled.
D_S_TOV	5,000 msec
E_D_TOV	2,000 msec
R_A_TOV	10,000 msec

Parameters	Default
Interop mode	Disabled

Configuring Timer 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 DCNM-SAN, expand **Switches > FC Services** and then select **Timers & Policies** in the Physical Attributes pane. You see the timers for multiple switches in the Information pane. Click the **Change Timeouts** button to configure the timeout values.

To configure timers in Device Manager, click **FC > Advanced > Timers/Policies**. You see the timers for a single switch in the dialog box.

Task Flow for Configuring Time Across All VSANs

Follow these steps to configure time across all VSANs:

Procedure

- Step 1** Configure the timer per-VSAN.
- Step 2** Enable the fctimer distribution.
- Step 3** Make the required configuration changes and committ the fctimer changes.
- Step 4** Discard the changes if you choose to discard the configuration changes.

Configuring Timer Per-VSAN

You can also issue the fctimer for a specified VSAN to configure different TOV values for VSANs with special links like FC or IP tunnels. You can configure different E_D_TOV, R_A_TOV, and D_S_TOV values for individual VSANs. Active VSANs are suspended and activated when their timer values are changed.



Caution

You cannot perform a nondisruptive downgrade to any earlier version that does not support per-VSAN FC timers.



Note This configuration must be propagated to all switches in the fabric—be sure to configure the same value in all switches in the fabric.

If a switch is downgraded to Cisco MDS SAN-OS Release 1.2 or 1.1 after the timer is configured for a VSAN, an error message is issued to warn against strict incompatibilities. Refer to the *Cisco MDS 9000 Family Troubleshooting Guide*.

To configure per-VSAN Fiber Channel timers using Device Manager, follow these steps:

Procedure

-
- Step 1** Click FC > Advanced > VSAN Timers.
You see the VSANs Timer dialog box.
 - Step 2** Fill in the timer values that you want to configure.
 - Step 3** Click Apply to save these changes.
-

Enabling fctimer Distribution

To enable and distribute fctimer configuration changes using Device Manager, follow these steps:

Procedure

-
- Step 1** Choose FC > Advanced > VSAN Timers.
You see the VSANs Timer dialog box.
 - Step 2** Fill in the timer values that you want to configure.
 - Step 3** Click Apply to save these changes.
 - Step 4** Select commit from the CFS drop-down menu to distribute these changes or select abort from the CFS drop-down menu to discard any unsaved changes.
-

Configuring a Secondary MAC Address

To allocate secondary MAC addresses using Device Manager, follow these steps:

Procedure

-
- Step 1** Choose FC > Advanced > WWN Manager.
You see the list of allocated WWNs.
 - Step 2** Supply the BaseMacAddress and MacAddressRange fields.

- Step 3** Click Apply to save these changes, or click Close to discard any unsaved changes.
-

Configuring Interop Mode 1

To configure interop mode 1 for a VSAN, follow these steps:

Procedure

- Step 1** Choose VSAN > VSAN Attributes from the Logical Domains pane.
- Step 2** Select Interop-1 from the Interop drop-down menu.
- Step 3** Click Apply Changes to save this interop mode.
- Step 4** Expand VSAN and then select Domain Manager from the Logical Domains pane.
- You see the Domain Manager configuration in the Information pane.
- Step 5** Set the Domain ID in the range of 97 (0x61) through 127 (0x7F).
- Click the **Configuration** tab.
 - Click in the Configure Domain ID column under the Configuration tab.
 - Click the **Running** tab and check that the change has been made.
- Note** This is a limitation imposed by the McData switches.
- Note** When changing the domain ID, the FC IDs assigned to N ports also change.
- Step 6** Change the Fibre Channel timers (if they have been changed from the system defaults).
- Note** The Cisco MDS 9000, Brocade, and McData FC error detect (ED_TOV) and resource allocation (RA_TOV) timers default to the same values. They can be changed if needed. The RA_TOV default is 10 seconds, and the ED_TOV default is 2 seconds. Per the FC-SW2 standard, these values must be the same on each switch within the fabric.
- Expand Switches > FC Services, and then select 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 VSAN > Domain Manager > Configuration and select disruptive or nonDisruptive in the Restart column to restart the domain.
-

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.

Verifying Interoperating Status

To verify the interoperability status of any switch in the Cisco MDS 9000 Family using DCNM for SAN, follow these steps:

Procedure

- Step 1** Choose Switches in the Physical Attributes pane and check the release number in the Information pane to verify the Cisco NX-OS release.
 - Step 2** Expand Switches > Interfaces, and then select FC Physical to verify the interface modes for each switch.
 - Step 3** Expand Fabricxx in the Logical Domains pane and then select All VSANs to verify the interop mode for all VSANs.
 - Step 4** Expand Fabricxx > All VSANs and then select Domain Manager to verify the domain IDs, local, and principal sWWNs for all VSANs.
 - Step 5** Using Device Manager, choose FC > Name Server to verify the name server information.
You see the Name Server dialog box.
 - Step 6** Click Close to close the dialog box.
-



CHAPTER 28

Configuring Users and Common Role

- [Configuring Users and Common Role, on page 599](#)

Configuring Users and Common Role

The CLI and SNMP use common roles in all switches in the Cisco MDS 9000 Family. You can use the CLI to modify a role that was created using SNMP and vice versa.

Users, passwords, and roles for all CLI and SNMP users are the same. A user configured through the CLI can access the switch using SNMP (for example, DCNM for SAN (DCNM-SAN or Device Manager) and vice versa.

This chapter includes the following topics:

Information About Role-Based Authorization

Switches in the Cisco MDS 9000 Family perform authentication based on roles. Role-based authorization limits access to switch operations by assigning users to roles. This kind of authentication restricts you to management operations based on the roles to which you have been added.

When you execute a command, perform command completion, or obtain context-sensitive help, the switch software allows the operation to progress if you have permission to access that command.

This section includes the following topics:

About Roles

Each role can contain multiple users and each user can be part of multiple roles. For example, if role1 users are only allowed access to configuration commands, and role2 users are only allowed access to **debug** commands, then if Joe belongs to both role1 and role2, he can access configuration as well as **debug** commands.



Note

If you belong to multiple roles, you can execute a union of all the commands permitted by these roles. Access to a command takes priority over being denied access to a command. For example, suppose you belong to a TechDocs group and you were denied access to configuration commands. However, you also belong to the engineering group and have access to configuration commands. In this case, you will have access to configuration commands.

**Tip**

Any role, when created, does not allow access to the required commands immediately. The administrator must configure appropriate rules for each role to allow access to the required commands.

Rules and Features for Each Role

Up to 16 rules can be configured for each role. These rules reflect what CLI commands are allowed. The user-specified rule number determines the order in which the rules are applied. For example, rule 1 is applied before rule 2, which is applied before rule 3, and so on. A user not belonging to the network-admin role cannot perform commands related to roles.

For example, if user A is permitted to perform all **show** CLI commands, user A cannot view the output of the **show role** CLI command if user A does not belong to the network-admin role.

A **rule** specifies operations that can be performed by a specific role. Each rule consists of a rule number, a rule type (permit or deny), a CLI command type (for example, **config**, **clear**, **show**, **exec**, **debug**), and an optional feature name (for example, FSPF, zone, VSAN, fcping, or interface).

**Note**

In this case, **exec** CLI commands refer to all commands in the EXEC mode that are not included in the **show**, **debug**, and **clear** CLI command categories.

Rule Changes Between SAN-OS Release 3.3(1c) and NX-OS Release 4.2(1a) Affect Role Behavior

The rules that can be configured for roles were modified between SAN-OS Release 3.3(1c) and NX-OS Release 4.2(1a). As a result, roles do not behave as expected following an upgrade from SAN-OS Release 3.3(1c) to NX-OS Release 4.2(1a). Manual configuration changes are required to restore the desired behavior.

Rule 4 and Rule 3: after the upgrade, exec and feature are removed. Change rule 4 and rule 3 as follows:

SAN-OS Release 3.3(1c) Rule	NX-OS Release 4.2(1a), Set the Rule to:
rule 4 permit exec feature debug	rule 4 permit debug
rule 3 permit exec feature clear	rule 3 permit clear

Rule 2: after the upgrade, exec feature license is obsolete.

SAN-OS Release 3.3(1c) Rule	NX-OS Release 4.2(1a) Rule
rule 2 permit exec feature debug	Not available in Release 4.2(1).

Rule 9, Rule 8, and Rule 7: after the upgrade, you need to have the feature enabled to configure it. In SAN-OS Release 3.3(1c), you could configure a feature without enabling it.

SAN-OS Release 3.3(1c) Rule	NX-OS Release 4.2(1a), to Preserve the Rule:
rule 9 deny config feature telnet	Not available in Release 4.2(1) and cannot be used.
rule 8 deny config feature tacacs-server	During the upgrade, enable the feature to preserve the rule; otherwise, the rule disappears.

SAN-OS Release 3.3(1c) Rule	NX-OS Release 4.2(1a), to Preserve the Rule:
rule 7 deny config feature tacacs+	During the upgrade, enable the feature to preserve the rule; otherwise, the rule disappears.

About the VSAN Policy

Configuring the VSAN policy requires the ENTERPRISE_PKG license (For more information, see Cisco MDS 9000 Family NX-OS Licensing Guide).

You can configure a role so that it only allows tasks to be performed for a selected set of VSANs. By default, the VSAN policy for any role is permit, which allows tasks to be performed for all VSANs. You can configure a role that only allows tasks to be performed for a selected set of VSANs. To selectively allow VSANs for a role, set the VSAN policy to deny, and then set the configuration to permit or the appropriate VSANs.



Note Users configured in roles where the VSAN policy is set to deny cannot modify the configuration for E ports. They can only modify the configuration for F or FL ports (depending on whether the configured rules allow such configuration to be made). This is to prevent such users from modifying configurations that may impact the core topology of the fabric.



Tip Roles can be used to create VSAN administrators. Depending on the configured rules, these VSAN administrators can configure MDS features (for example, zone, fcdomain, or VSAN properties) for their VSANs without affecting other VSANs. Also, if the role permits operations in multiple VSANs, then the VSAN administrators can change VSAN membership of F or FL ports among these VSANs.

Users belonging to roles in which the VSAN policy is set to deny are referred to as VSAN-restricted users.

Role Distributions

Role-based configurations use the Cisco Fabric Services (CFS) infrastructure to enable efficient database management, and to provide a single point of configuration for the entire fabric.

The following configurations are distributed:

- Role names and descriptions
- List of rules for the roles
- VSAN policy and the list of permitted VSANs

About Role Databases

Role-based configurations use two databases to accept and implement configurations.

- Configuration database—The running database currently enforced by the fabric.
- Pending database—Your subsequent configuration changes are stored in the pending database. If you modify the configuration, you need to commit or discard the pending database changes to the configuration database. The fabric remains locked during this period. Changes to the pending database are not reflected in the configuration database until you commit the changes.

Locking the Fabric

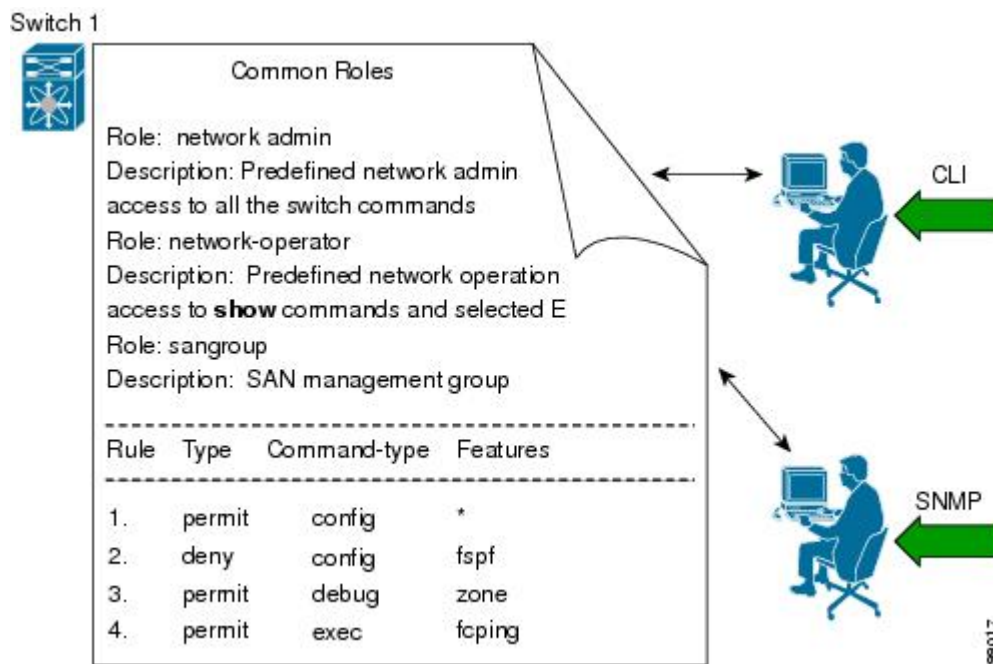
The first action that modifies the database creates the pending database and locks the feature in the entire fabric. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database along with the first change.

About Common Roles

The CLI and SNMP in all switches in the Cisco MDS 9000 Family use common roles. You can use SNMP to modify a role that was created using the CLI and vice versa (see [Figure 76: Common Roles](#), on page 602).

Figure 76: Common Roles



Each role in SNMP is the same as a role created or modified through the CLI (see the [Information About Role-Based Authorization](#), on page 599).

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

You can create new roles or modify existing roles using SNMP or the CLI.

- SNMP—Use the CISCO-COMMON-ROLES-MIB to configure or modify roles. Refer to the *Cisco MDS 9000 Family MIB Quick Reference*.
- CLI—Use the **role name** command.

Mapping of CLI Operations to SNMP

SNMP has only three possible operations: GET, SET, and NOTIFY. The CLI has five possible operations: DEBUG, SHOW, CONFIG, CLEAR, and EXEC.



Note NOTIFY does not have any restrictions like the syslog messages in the CLI.

Table 72: CLI Operation to SNMP Operation Mapping , on page 603 explains how the CLI operations are mapped to the SNMP operations.

Table 72: CLI Operation to SNMP Operation Mapping

CLI Operation	SNMP Operation
DEBUG	Ignored
SHOW	GET
CONFIG	SET
CLEAR	SET
EXEC	SET

Creating Users Guidelines

The passphrase specified in the **snmp-server user** option and the password specified **username** option are synchronized.

By default, the user account does not expire unless you explicitly configure it to expire. The **expire** option determines the date on which the user account is disabled. The date is specified in the YYYY-MM-DD format.

When creating users, note the following guidelines:

- You can configure up to a maximum of 256 users on a switch.
- The following words are reserved and cannot be used to configure users: bin, daemon, adm, lp, sync, shutdown, halt, mail, news, uucp, operator, games, gopher, ftp, nobody, nsd, mailnull, rpc, rpcuser, xfs, gdm, mtsuser, ftpuser, man, and sys.
- User passwords are not displayed in the switch configuration file.
- If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password as shown in the sample configuration. Passwords are case-sensitive. “admin” is no longer the default password for any Cisco MDS 9000 Family switch. You must explicitly configure a strong password.
- To issue commands with the **internal** keyword for troubleshooting purposes, you must have an account that is a member of the network-admin group.



Caution

Cisco MDS NX-OS supports user names that are created with alphanumeric characters or specific special characters (+ [plus], = [equal], _ [underscore], - [hyphen], \ [backslash], and . [period]) whether created remotely (using TACACS+ or RADIUS) or locally. Local user names cannot be created with any special characters (apart from those specified). If a non-supported special character user name exists on an AAA server, and is entered during login, then the user is denied access.

Characteristics of Strong Passwords

A strong password has the following characteristics:

- At least eight characters long
- Does not contain many consecutive characters (such as “abcd”)
- Does not contain many repeating characters (such as “aaabbb”)
- Does not contain dictionary words
- Does not contain proper names
- Contains both upper- and lower-case characters
- Contains numbers

The following are examples of strong passwords:

- If2CoM18
- 2004AsdfLkj30
- Cb1955S21

About SSH

SSH provides secure communications to the Cisco NX-OS CLI. You can use SSH keys for the following SSH options:

- SSH2 using RSA
- SSH2 using DSA

Boot Mode SSH

Due to the increasing emphasis on security and security-related issues, the ssh command in this release runs in the Boot mode. SSH is a preferred and more secure method of data exchange over the network because it communicates over the secure channel, and the data is encrypted before sending on the channel.

SSH Authentication Using Digital Certificates

SSH authentication on the Cisco MDS 9000 Family switches provide X.509 digital certificate support for host authentication. An X.509 digital certificate is a data item that vouches for the origin and integrity of a message. It contains encryption keys for secured communications and is “signed” by a trusted certification authority (CA) to verify the identity of the presenter. The X.509 digital certificate support provides either DSA or RSA algorithms for authentication.

The certificate infrastructure uses the first certificate that supports the Secure Socket Layer (SSL) and is returned by the security infrastructure, either through query or notification. Verification of certificates is successful if the certificates are from any of the trusted CAs.

You can configure your switch for either SSH authentication using an X.509 certificate or SSH authentication using a Public Key Certificate, but not both. If either of them is configured and the authentication fails, you will be prompted for a password.

For more information on CAs and digital certificates, see *Chapter 30, “Configuring Certificate Authorities and Digital Certificates.”*

Passwordless File copy and SSH

Secure Shell (SSH) public key authentication can be used to achieve password-free logins. SCP and SFTP uses SSH in the background, which enables these copy protocols to be used for a password-free copy with public key authentication. The NX-OS version only supports the SCP and STFP client functionality.

You can create an RSA and DSA identity that can be used for authentication with SSH. The identity consists of two parts: public and private keys. The public and the private keys are generated by the switch or can be generated externally and imported to the switch. For import purposes, the keys should be in OPENSSH format.

To use the key on a host machine hosting an SSH server, you must transfer the public key file to the machine and add the contents of it to the `authorized_keys` file in your SSH directory (for example, `$HOME/.ssh`) on the server. For the import and export of private keys, the key is protected by encryption. You are asked to enter the passphrase for the keys. If you enter a passphrase, the private key is protected by encryption. If you leave the password field blank, the key will not be encrypted.

If you need to copy the keys to another switch, you will have to export the keys out of the switch to a host machine, and then import the keys to other switches from that machine.

The key files are persistent across reload.

Guidelines and Limitations

Fabric merge does not modify the role database on a switch. If two fabrics merge, and the fabrics have different role databases, the software generates an alert message.

See the “*Merge Guidelines for RADIUS and TACACS+ Configurations*” section for detailed concepts.

- Verify that the role database is identical on all switches in the entire fabric.
- Be sure to edit the role database on any switch to the desired database and then commit it. This synchronizes the role databases on all the switches in the fabric.

Default Settings

[Table 73: Default Switch Security Settings](#), on page 605 lists the default settings for all switch security features in any switch.

Table 73: Default Switch Security Settings

Parameters	Default
Roles in Cisco MDS Switches	Network operator (network-operator)
AAA configuration services	Local
Authentication port	1812
Accounting port	1813
Preshared key communication	Clear text
RADIUS server time out	1 (one) second
RADIUS server retries	Once

Parameters	Default
TACACS+	Disabled
TACACS+ servers	None configured
TACACS+ server timeout	5 seconds
AAA server distribution	Disabled
VSAN policy for roles	Permit
User account	No expiry (unless configured)
Password	None
Password-strength	Enabled
Accounting log size	250 KB
SSH service	Enabled
Telnet service	Disabled

Configuring Users and Common Role

This section includes the following topics:

Configuring Roles and Profiles

To create an additional role or to modify the profile for an existing role, follow these steps:



Note

Only users belonging to the network-admin role can create roles.



Note

Device Manager automatically creates six roles that are required for Device Manager to display a view of a switch. These roles are **system**, **snmp**, **module**, **interface**, **hardware**, and **environment**.

Procedure

- Step 1** Expand Switches > Security and then select Users and Roles from the Physical Attributes pane.
- Step 2** Click the Roles tab in the Information pane.
- Step 3** Click Create Row to create a role in DCNM-SAN.
- Step 4** Select the switches on which to configure a role.
- Step 5** Enter the name of the role in the Name field.
- Step 6** Enter the description of the role in the Description field.

- Step 7** (Optional) Check the Enable check box to enable the VSAN scope and enter the list of VSANs in the Scope field to which you want to restrict this role.
- Step 8** Click Create to create the role.
-

Deleting Common Roles

To delete a common role using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand Switches > Security and then select Users and Roles from the Physical Attributes pane.
- Step 2** Click the Roles tab in the Information pane.
- Step 3** Click the role you want to delete.
- Step 4** Click Delete Row to delete the common role.
- Step 5** Click **Yes** to confirm the deletion or **No** to cancel it.
-

Modifying Rules

To modify the rules for an existing role using Device Manager, follow these steps:

Procedure

- Step 1** Choose Security > Roles.
- Step 2** Click the role for which you want to edit the rules.
- Step 3** Click Rules to view the rules for the role.
- You see the Edit Role Rules dialog box.
- Step 4** Edit the rules you want to enable or disable for the common role.
- Step 5** Click Apply to apply the new rules.
-

What to do next

Rule 1 is applied first, which permits, for example, sangroup users access to all **config** CLI commands. Rule 2 is applied next, denying FSPF configuration to sangroup users. As a result, sangroup users can perform all other **config** CLI commands, except the **fspf** CLI configuration commands.



Note The order of rule placement is important. If you had swapped these two rules and issued the **deny config feature fsfp** rule first and issued the **permit config** rule next, you would be allowing all sangroup users to perform all configuration commands because the second rule globally overrode the first rule.

Modifying the VSAN Policy

To modify the VSAN policy for an existing role, follow these steps:

Procedure

- Step 1** Expand Switches > Security and then select Users and Roles from the Physical Attributes pane.
 - Step 2** Click the Roles tab in the Information pane.
 - Step 3** Check the **Scope** Enable check box if you want to enable the VSAN scope and restrict this role to a subset of VSANs.
 - Step 4** Enter the list of VSANs in the Scope VSAN Id List field that you want to restrict this role to.
 - Step 5** Click Apply Changes to save these changes.
-

Committing Role-Based Configuration Changes

If you commit the changes made to the pending database, the configuration is committed to all the switches in the fabric. On a successful commit, the configuration change is applied throughout the fabric and the lock is released. The configuration database now contains the committed changes and the pending database is now cleared.

To commit role-based configuration changes, follow these steps:

Procedure

- Step 1** Expand Switches > Security and then select Users and Roles in the Physical Attributes pane.
 - Step 2** Click the Roles CFS tab in the Information pane.
 - Step 3** Set the Global drop-down menu to enable to enable CFS.
 - Step 4** Click the Apply Changes icon to save this change.
 - Step 5** Set the Config Action drop-down menu to commit to commit the roles using CFS.
 - Step 6** Click the Apply Changes icon to save this change.
-

Discarding Role-Based Configuration Changes

If you discard (abort) the changes made to the pending database, the configuration database remains unaffected and the lock is released.

To discard role-based configuration changes, follow these steps:

Procedure

- Step 1** Expand Switches > Security and then select Users and Roles in the Physical Attributes pane.
- Step 2** Click the Roles **CFS** tab in the Information pane.
- Step 3** Set the Config Action drop-down menu to abort to discard any uncommitted changes.

- Step 4** Click the Apply Changes icon to save this change.
-

Enabling Role-Based Configuration Distribution

To enable role-based configuration distribution, follow these steps:

Procedure

- Step 1** Expand Switches > Security and then select Users and Roles in the Physical Attributes pane.
- Step 2** Click the Roles **CFS** tab in the Information pane.
- Step 3** Set the Global drop-down menu to enable to enable CFS distribution.
- Step 4** Click the Apply Changes icon to save this change.
-

Clearing Sessions

To forcibly clear the existing role session in the fabric using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand Switches > Security and then select Users and Roles in the Physical Attributes pane.
- Step 2** Click the Roles **CFS** tab in the Information pane.
- Step 3** Set the Config Action drop-down menu to clear to clear the pending database.
- Step 4** Click the Apply Changes icon to save this change.
-

What to do next



Caution Any changes in the pending database are lost when you clear a session.

Configuring Users

Before configuring users, make sure that you have configured roles to associate with the users that you are creating.



Note As of Cisco SAN-OS Release 3.1(2b), DCNM-SAN automatically checks whether encryption is enabled, which allows you to create users.

To configure a new user or to modify the profile of an existing user, follow these steps:

Procedure

Step 1 Expand Switches > Security and then select Users and Roles from the Physical Attributes pane.

Step 2 Click the Users tab in the Information pane to see a list of users.

Step 3 Click the Create Row icon.

You see the Users - Create dialog box as shown in [Figure 77: Users - Create Dialog Box, on page 610](#).

Figure 77: Users - Create Dialog Box

Step 4 (Optional) Alter the Switches check boxes to specify one or more switches.

Step 5 Enter the user name in the New User field.

Step 6 Enter the password for the user.

Step 7 Check the roles that you want to associate with this user.

See the [Rules and Features for Each Role, on page 600](#).

Step 8 Select the appropriate option for the type of authentication protocol used. The default value is MD5.

Step 9 Select the appropriate option for the type of privacy protocol used. The default value is DES.

Step 10 (Optional) Enter the expiry date for this user.

Step 11 (Optional) Enter the SSH Key filename.

Step 12 Click Create to create the entry.

Deleting a User

To delete a user using DCNM-SAN, follow these steps:

Procedure

-
- | | |
|---------------|-----------------------------------------------------------------------------------------------|
| Step 1 | Expand Switches > Security and then select Users and Roles from the Physical Attributes pane. |
| Step 2 | Click the Users tab in the Information pane to see a list of users. |
| Step 3 | Click the name of the user you want to delete. |
| Step 4 | Click Delete Row to delete the selected user. |
| Step 5 | Click Apply Changes to save this change. |
-

Configuring SSH Services

A secure SSH connection with an RSA key is available as a default on all Cisco MDS 9000 Family switches. If you require a secure SSH connection with a DSA key, you need to disable the default SSH connection, Generate a DSA key and then enable the SSH connection (see the [Generating the SSH Server Key Pair](#), on page 611).

Use the **ssh key** command to generate a server key.



Caution	If you are logging in to a switch through SSH and you have issued the aaa authentication login default none command, you must enter one or more key strokes to log in. If you press the Enter key without entering at least one keystroke, your log in will be rejected.
----------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

This section includes the following topics:

Generating the SSH Server Key Pair

Ensure that you have an SSH server key pair with the appropriate version before enabling the SSH service. Generate the SSH server key pair according to the SSH client version used. The number of bits specified for each key pair ranges from 768 to 2048.

The SSH service accepts two types of key pairs for use by SSH version 2.

- The **dsa** option generates the DSA key pair for the SSH version 2 protocol.
- The **rsa** option generates the RSA keypair for the SSH version 2 protocol.



Caution	If you delete all of the SSH keys, you cannot start a new SSH session.
----------------	------------------------------------------------------------------------

To generate the SSH server key pair, follow these steps:

Overwriting a Generated Key Pair


Procedure

Step 1 Expand Switches > Security and then select SSH and Telnet.

Step 2 Click the Create Row icon.

You see the SSH and Telnet Key - Create dialog box (see [Figure 78: SSH and Telnet - Create Dialog Box](#), on page 612).

Figure 78: SSH and Telnet - Create Dialog Box



Step 3 Check the switches you want to assign to this SSH key pair.

Step 4 Choose the key pair option type from the listed Protocols. The listed protocols are SSH1, SSH2(rsa), and SSH2(dsa).

Step 5 Set the number of bits that will be used to generate the key pairs in the NumBits drop-down menu.

Step 6 Click Create to generate these keys.

Note 1856 DSA NumberKeys are not supported by switches that running Cisco MDS NX-OS software version 4.1(1) and later.

Overwriting a Generated Key Pair

If the SSH key pair option is already generated for the required version, you can force the switch to overwrite the previously generated key pair.

To overwrite the previously generated key pair, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	Expand Switches > Security and then select SSH and Telnet.	You see the configuration in the Information pane.

	Command or Action	Purpose
Step 3	Highlight the key that you want to overwrite and click Delete Row.	
Step 4	Click the Apply Changes icon to save these changes.	
Step 5	Click the Create Row icon.	You see the SSH and Telnet Key - Create dialog box.
Step 6	Check the switches you want to assign this SSH key pair.	
Step 7	Choose the key pair option type from the Protocols radio buttons.	
Step 8	Set the number of bits that will be used to generate the key pairs in the NumBits drop-down menu.	
Step 9	Click Create to generate these keys.	

Enabling SSH or Telnet Service


By default, the SSH service is enabled with the RSA key.

To enable or disable the SSH or Telnet service, follow these steps:

Procedure

- Step 1** Expand Switches > Security and then select SSH and Telnet.
- Step 2** Select the Control tab and check an SSH check box or **Telnet** check box for each switch (see [Figure 79: Control Tab under SSH and Telnet, on page 613](#)).

Figure 79: Control Tab under SSH and Telnet



- Step 3** Click the Apply Changes icon to save this change.

Changing Administrator Password Using DCNM-SAN

To change the administrator password in DCNM-SAN, follow these steps:

Procedure

Step 1 Click the Open tab in the control panel.

Step 2 Choose the password field to change the password for an already existing user for the fabric.

Step 3 Click Open to open the fabric.

Note New password will be saved after the fabric is open. The user name and password fields are editable in the Fabric tab only after you unmanage the fabric.

Recovering the Administrator Password

You can recover the administrator password using one of two methods:

- From the CLI with a user name that has network-admin privileges.
- Power cycling the switch.



Note To recover an administrator's password, refer to the Cisco MDS 9000 Family NX-OS Security Configuration Guide.

The following topics included in this section:

Displaying Role-Based Information

The rules are displayed by rule number and are based on each role. All roles are displayed if the role name is not specified.

To view rules for a role using Device Manager, follow these steps:

Procedure

Step 1 Click Security > Roles.

You see the Roles dialog box.

Step 2 Select a role name and click Rules.

You see the Rules dialog box.

Step 3 Click **Summary** to get a summarized view of the rules configured for this role.

Displaying Roles When Distribution is Enabled

Procedure

- Step 1** Expand Switches > Security and then select Users and Roles in the Physical Attributes pane.
- Step 2** Click the Users tab in the Information pane(see [Figure 80: Roles CFS Tab, on page 615](#)).

Figure 80: Roles CFS Tab

SAN/Fabric sw-189/Switches/Security/Users

Roles CFS

Roles

Users

Communities

Global

Switch	Feature Admin	Feature Oper	Global State	Config Action	Last Command	Last Result	Lock Owner Switch	Lock Owner User Name	Merge Status	Master	Scope
Y-172.22.31.184	<div>noSelection</div>	<div>disabled</div>	<div>disable</div>	<div>noSelection</div>						<input type="checkbox"/>	rfFabric ipNetwork
y-188	<div>noSelection</div>	<div>disabled</div>	<div>enable</div>	<div>noSelection</div>					<div>failure...</div>	<input type="checkbox"/>	rfFabric ipNetwork
y-185	<div>noSelection</div>	<div>enabled</div>	<div>enable</div>	<div>noSelection</div>					<div>failure...</div>	<input checked="" type="checkbox"/>	rfFabric ipNetwork
y-190	<div>noSelection</div>	<div>enabled</div>	<div>enable</div>	<div>noSelection</div>					<div>failure...</div>	<input type="checkbox"/>	rfFabric ipNetwork
c-186	<div>noSelection</div>	<div>enabled</div>	<div>enable</div>	<div>noSelection</div>					<div>failure...</div>	<input type="checkbox"/>	rfFabric ipNetwork
sw-189	<div>noSelection</div>	<div>disabled</div>	<div>disable</div>	<div>noSelection</div>					<div>failure...</div>	<input type="checkbox"/>	rfFabric ipNetwork

240482

- Step 3** Set the Config View As drop-down value to pending to view the pending database or set the Config View as drop-down menu to running to view the running database.
- Step 4** Click Apply Changes to save this change.

Displaying User Account Information

To display information about configured user accounts using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand Security and then select Users and Roles in the Physical Attributes pane.
- Step 2** Click the Users tab.

You see the list of SNMP users shown in [Figure 81: Users Listed Under the Users Tab, on page 615](#) in the Information pane.

Figure 81: Users Listed Under the Users Tab

/SAN/fabric.172.22.46.220/VSAN0001/Switches/Security/Users and Roles									
Roles CFS	Roles	Users	Communities	Privacy					
Switch	User	Role	Password (not echoed)	Digest	Encryption	ExpiryDate (eg. yyyy/mm/dd-hh:mm:ss)	SSH Key File Configured	SSH Key File (bootflash:[volatile:]) (not echoed)	Creation 1
sw172-22-46-174	admin	network-admin		MD5	DES		False	localCredr	
sw172-22-46-174	mcchin	network-admin, network-operator		NoAuth	NoPriv		False	localCredr	
sw172-22-46-174	mdsusr	network-admin, network-operator		NoAuth	NoPriv		False	localCredr	
sw172-22-46-174	shausr	network-admin		NoAuth	NoPriv		False	localCredr	
sw172-22-46-220	admin	network-admin		MD5	DES		False	localCredr	
sw172-22-46-220	assusr	network-admin, network-operator		NoAuth	NoPriv		False	localCredr	
sw172-22-46-220	madmin	network-admin, network-operator		NoAuth	NoPriv		False	localCredr	
sw172-22-46-220	mcchin	network-admin, network-operator		MD5	DES		False	localCredr	
sw172-22-46-220	mdsusr	network-admin, network-operator		NoAuth	NoPriv		False	localCredr	
sw172-22-46-220	newusr	network-admin, network-operator		NoAuth	NoPriv		False	localCredr	
sw172-22-46-220	shausr	network-admin, network-operator		NoAuth	NoPriv		False	localCredr	
sw172-22-46-220	imgtusr	network-admin, network-operator		NoAuth	NoPriv		False	localCredr	

Field Descriptions for Users and Common Role

Common Roles



Note Common roles is not available in displayFCoE mode (use security roles).

Field	Description
Description	Description of the common role.
Enable	This specifies whether the common role has a VSAN restriction or not.
List	List of VSANs user is restricted to.

Feature History for Users and Common Role

[Table 74: Feature History for FIPS](#) , on page 616 lists the release history for this feature. Only features that were introduced or modified in 5.x or a later release appear in the table.

Table 74: Feature History for FIPS

Feature Name	Releases	Feature Information
Changes to SSH	5.0(1a)	Boot Mode SSH, Passwordfree File copy, and SSH.
Role Distributions	5.0(1a)	Enabling role-based configuration distribution.
Creating Users Guidelines	5.0(1a)	Caution has been changed.



CHAPTER 29

Configuring Security Features on External AAA Server

- [Configuring Security Features on an External AAA Server, on page 617](#)

Configuring Security Features on an External AAA Server

The authentication, authorization, and accounting (AAA) feature verifies the identity of, grants access to, and tracks the actions of users managing a switch. All Cisco MDS 9000 Family switches use Remote Access Dial-In User Service (RADIUS) or Terminal Access Controller Access Control device Plus (TACACS+) protocols to provide solutions using remote AAA servers.

Based on the user ID and password combination provided, switches perform local authentication or authorization using the local database or remote authentication or authorization using a AAA server. A preshared secret key provides security for communication between the switch and AAA servers. This secret key can be configured for all AAA servers or for only a specific AAA server. This security feature provides a central management capability for AAA servers.

This chapter includes the following topics:

Information About Switch Management Security

Management security in any switch in the Cisco MDS 9000 Family provides security to all management access methods, including the command-line interface (CLI) or Simple Network Management Protocol (SNMP).

This section includes the following topics:

Security Options

You can access the CLI using the console (serial connection), Telnet, or Secure Shell (SSH). You can access DCNM-SAN using TCP/UDP SNMP or HTTP traffic. For each management path (console, Telnet, and SSH), you can configure one or more of the following security control options: local, remote (RADIUS or TACACS+), or none.

- Remote security control
 - Using RADIUS

See the [Configuring the RADIUS, TACACS+, and LDAP Server, on page 635](#)

- Using TACACS+

See the [Configuring the RADIUS, TACACS+, and LDAP Server, on page 635](#)

- Local security control.

See the [Local AAA Services, on page 631](#).

These security features can also be configured for the following scenarios:

- iSCSI authentication

See the IP Services Configuration Guide, Cisco DCNM for SAN.

- Fibre Channel Security Protocol (FC-SP) authentication

See *Configuring FC-SP and DHCHAP*.

SNMP Security Options

The SNMP agent supports security features for SNMPv1, SNMPv2c, and SNMPv3. Normal SNMP security features apply to all applications that use SNMP (for example, Cisco MDS 9000 DCNM for SAN).

SNMP security options also apply to DCNM for SAN and Device Manager.

See the Cisco MDS 9000 NX-OS Family System Management Configuration Guide for more information on the SNMP security options.

Refer to the *Cisco DCNM Fundamentals Guide* for information on DCNM for SAN and Device Manager.

Switch AAA Functionalities

Using the CLI or DCNM for SAN (DCNM-SAN), or an SNMP application, you can configure AAA switch functionalities on any switch in the Cisco MDS 9000 Family.

This section includes the following topics:

Authentication

Authentication is the process of verifying the identity of the person or device accessing the switch. This identity verification is based on the user ID and password combination provided by the entity trying to access the switch. Cisco MDS 9000 Family switches allow you to perform local authentication (using the local lookup database) or remote authentication (using one or more RADIUS or TACACS+ servers).



Note

When you log in to a Cisco MDS switch successfully using DCNM-SAN or Device Manager through Telnet or SSH and if that switch is configured for AAA server-based authentication, a temporary SNMP user entry is automatically created with an expiry time of one day. The switch authenticates the SNMPv3 protocol data units (PDUs) with your Telnet or SSH login name as the SNMPv3 user. The management station can temporarily use the Telnet or SSH login name as the SNMPv3 **auth** and **priv** passphrase. This temporary SNMP login is only allowed if you have one or more active MDS shell sessions. If you do not have an active session at any given time, your login is deleted and you will not be allowed to perform SNMPv3 operations.



Note DCNM-SAN does not support AAA passwords with trailing white space, for example “passwordA.”

Authorization

The following authorization roles exist in all Cisco MDS switches:

- Network operator (network-operator): Has permission to view the configuration only. The operator cannot make any configuration changes.
- Network administrator (network-admin): Has permission to execute all commands and make configuration changes. The administrator can also create and customize up to 64 additional roles.
- Default-role: Has permission to use the GUI (DCNM-SAN and Device Manager). This access is automatically granted to all users for accessing the GUI.

These roles cannot be changed or deleted. You can create additional roles and configure the following options:

- Configure role-based authorization by assigning user roles locally or using remote AAA servers.
- Configure user profiles on a remote AAA server to contain role information. This role information is automatically downloaded and used when the user is authenticated through the remote AAA server.



Note If a user belongs only to one of the newly created roles and that role is subsequently deleted, then the user immediately defaults to the network-operator role.

Accounting

The accounting feature tracks and maintains a log of every management configuration used to access the switch. This information can be used to generate reports for troubleshooting and auditing purposes. Accounting logs can be stored locally or sent to remote AAA servers.

Remote AAA Services

Remote AAA services provided through RADIUS and TACACS+ protocols have the following advantages over local AAA services:

- User password lists for each switch in the fabric can be managed more easily.
- AAA servers are already deployed widely across enterprises and can be easily adopted.
- The accounting log for all switches in the fabric can be centrally managed.
- User role mapping for each switch in the fabric can be managed more easily.

Server Groups

You can specify remote AAA servers for authentication, authorization, and accounting using server groups. A server group is a set of remote AAA servers implementing the same AAA protocol. The purpose of a server group is to provide for failover servers in case a remote AAA server fails to respond. If the first remote server in the group fails to respond, the next remote server in the group is tried until one of the servers sends a response. If all the AAA servers in the server group fail to respond, then that server group option is considered a failure. If required, you can specify multiple server groups. If the Cisco MDS switch encounters errors from the servers in the first group, it tries the servers in the next server group.

AAA Service Configuration Options

AAA configuration in Cisco MDS 9000 Family switches is service based. You can have separate AAA configurations for the following services:

- Telnet or SSH login (DCNM-SAN and Device Manager login)
- Console login
- iSCSI authentication (see the Cisco MDS 9000 Family NX-OS IP Services Configuration Guide, Cisco DCNM for SAN).
- FC-SP authentication (see *Configuring FC-SP and DHCHAP*).
- Accounting

In general, server group, local, and none are the three options that can be specified for any service in an AAA configuration. Each option is tried in the order specified. If all the options fail, local is tried.



Caution

Cisco MDS NX-OS supports user names that are created with alphanumeric characters or specific special characters (+ [plus], = [equal], _ [underscore], - [hyphen], \ [backslash], and . [period]) whether created remotely (using TACACS+ or RADIUS) or locally, provided the user name starts with an alphabetical character. Local user names cannot be created with all numbers or with any special characters (apart from those specified). If a numeric-only user name or a non-supported special character user name exists on an AAA server, and is entered during login, then the user is denied access.



Note

Even if local is not specified as one of the options, it is tried by default if all AAA servers configured for authentication are unreachable. User has the flexibility to disable this fallback (See section [Configuring Fallback Mechanism for Authentication, on page 636](#)).

When RADIUS times out, local login is attempted depending on the fallback configuration. For this local login to be successful, a local account for the user with the same password should exist, and the RADIUS timeout and retries should take less than 40 seconds. The user is authenticated if the username and password exist in the local authentication configuration.

[Table 75: AAA Service Configuration Commands , on page 620](#) provides the related CLI command for each AAA service configuration option.

Table 75: AAA Service Configuration Commands

AAA Service Configuration Option	Related Command
Telnet or SSH login (Cisco DCNM-SAN and Device Manager login)	aaa authentication login default
Console login	aaa authentication login console
iSCSI authentication	aaa authentication iscsi default
FC-SP authentication	aaa authentication dhchap default
Accounting	aaa accounting default

Error-Enabled Status

When you log in, the login is processed by rolling over to local user database if the remote AAA servers do not respond. In this situation, the following message is displayed on your screen if you have enabled the error-enabled feature:

```
Remote AAA servers unreachable; local authentication done.
```

To enable this message display, use the **aaa authentication login error-enable** command.

To disable this message display, use the **no aaa authentication login error-enable** command.

To view the current display status, use the **show aaa authentication login error-enable** command (see).

[Displays AAA Authentication Login Information, on page 621](#)

Displays AAA Authentication Login Information

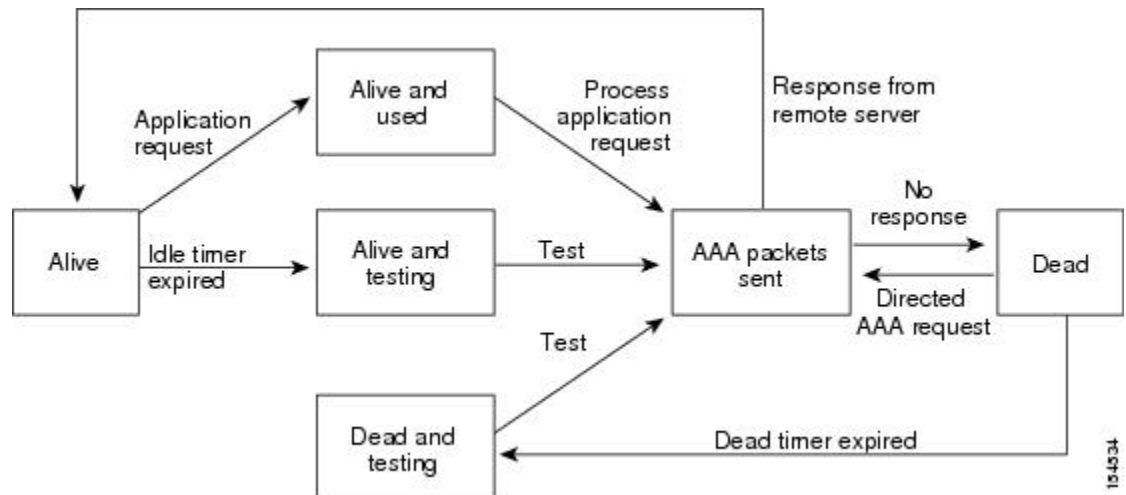
```
switch# show aaa authentication login error-enable
enabled
```

AAA Server Monitoring

An unresponsive AAA server introduces a delay in the processing of AAA requests. An MDS switch can periodically monitor an AAA server to check whether it is responding (or alive) to save time in processing AAA requests. The MDS switch marks unresponsive AAA servers as dead and does not send AAA requests to any dead AAA servers. An MDS switch periodically monitors dead AAA servers and brings them to the alive state once they are responding. This monitoring process verifies that an AAA server is in a working state before real AAA requests are sent its way. Whenever an AAA server changes to the dead or alive state, an SNMP trap is generated and the MDS switch warns the administrator that a failure is taking place before it can impact performance.

See the following image for AAA server states.

Figure 82: AAA Server States





Note The monitoring interval for alive servers and dead servers is different and can be configured by the user. The AAA server monitoring is performed by sending a test authentication request to the AAA server.

The user name and password to be used in the test packet can be configured.

Authentication and Authorization Process

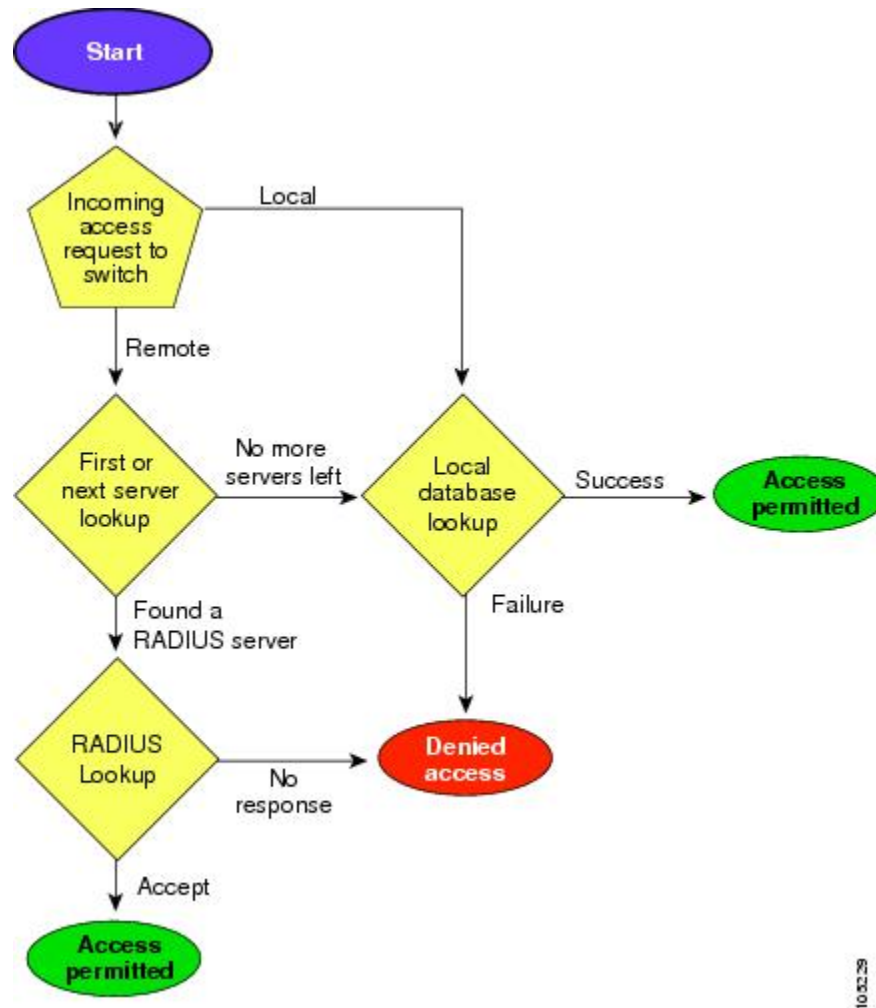
Authentication is the process of verifying the identity of the person managing the switch. This identity verification is based on the user ID and password combination provided by the person managing the switch. The Cisco MDS 9000 Family switches allow you to perform local authentication (using the lookup database) or remote authentication (using one or more RADIUS servers or TACACS+ servers).

Authorization provides access control. It is the process of assembling a set of attributes that describe what the user is authorized to perform. Based on the user ID and password combination, the user is authenticated and authorized to access the network as per the assigned role. You can configure parameters that can prevent unauthorized access by an user, provided the switches use the TACACS+ protocol.

AAA authorization is the process of assembling a set of attributes that describe what the user is authorized to perform. Authorization in the Cisco NX-OS software is provided by attributes that are downloaded from AAA servers. Remote security servers, such as RADIUS and TACACS+, authorize users for specific rights by associating attribute-value (AV) pairs, which define those rights with the appropriate user.

The following figure shows a flow chart of the authorization and authentication process.

Figure 83: Switch Authorization and Authentication Flow



Note No more server groups left = no response from any server in all server groups. No more servers left = no response from any server within this server group.

Global AAA Server Monitoring Parameters

The global AAA server monitoring parameters function as follows:

- When a new AAA server is configured it is monitored using the global test parameters, if defined.
- When global test parameters are added or modified, all the AAA servers, which do not have any test parameters configured, start getting monitored using the new global test parameters.
- When the server test parameters are removed for a server or when the idle-time is set to zero (default value) the server starts getting monitored using the global test parameters, if defined.
- If global test parameters are removed or global idle-time is set to zero, servers for which the server test parameters are present are not affected. However, monitoring stops for all other servers that were previously being monitored using global parameters.

- If the server monitoring fails with the user-specified server test parameters, the server monitoring does not fall back to global test parameters.

LDAP

The Lightweight Directory Access Protocol (LDAP) provides centralized validation of users attempting to gain access to a Cisco NX-OS device. LDAP services are maintained in a database on an LDAP daemon running, typically, on a UNIX or Windows NT workstation. You must have access to and must configure an LDAP server before the configured LDAP features on your Cisco NX-OS device are available.

LDAP provides for separate authentication and authorization facilities. LDAP allows for a single access control server (the LDAP daemon) to provide each service-authentication and authorization-independently. Each service can be tied into its own database to take advantage of other services available on that server or on the network, depending on the capabilities of the daemon.

The LDAP client/server protocol uses TCP (TCP port 389) for transport requirements. Cisco NX-OS devices provide centralized authentication using the LDAP protocol.



Note

If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

LDAP Authentication and Authorization

Clients establish a TCP connection and authentication session with an LDAP server through a simple bind (username and password). As part of the authorization process, the LDAP server searches its database to retrieve the user profile and other information.

You can configure the bind operation to first bind and then search, where authentication is performed first and authorization next, or to first search and then bind. The default method is to first search and then bind.

The advantage of searching first and binding later is that the distinguished name (DN) received in the search result can be used as the user DN during binding rather than forming a DN by prepending the username (cn attribute) with the baseDN. This method is especially helpful when the user DN is different from the username plus the baseDN. For the user bind, the bindDN is constructed as baseDN + append-with-baseDN, where append-with-baseDN has a default value of cn=\$userid.



Note

As an alternative to the bind method, you can establish LDAP authentication using the compare method, which compares the attribute values of a user entry at the server. For example, the user password attribute can be compared for authentication. The default password attribute type is userPassword.

About RADIUS Server Default Configuration

DCNM-SAN allows you to set up a default configuration that can be used for any RADIUS server that you configure the switch to communicate with. The default configuration includes:

- Encryption type
- Timeout value
- Number of retransmission attempts
- Allowing the user to specify a RADIUS server at login

About the Default RADIUS Server Encryption Type and Preshared Key

You need to configure the RADIUS preshared key to authenticate the switch to the RADIUS server. The length of the key is restricted to 64 characters and can include any printable ASCII characters (white spaces are not allowed). You can configure a global key to be used for all RADIUS server configurations on the switch.

You can override this global key assignment by explicitly using the **key** option when configuring an individual RADIUS server in the **radius-server host** command.

About RADIUS Servers

You can add up to 64 RADIUS servers. RADIUS keys are always stored in encrypted form in persistent storage. The running configuration also displays encrypted keys. When you configure a new RADIUS server, you can use the default configuration or modify any of the parameters to override the default RADIUS configuration.

Configuring the Test Idle Timer

The test idle timer specifies the interval during which a RADIUS server receives no requests before the MDS switch sends out a test packet.

**Note**

The default idle timer value is 0 minutes. When the idle time interval is 0 minutes, periodic RADIUS server monitoring is not performed.

Configuring Test User Name

You can configure a username and password for periodic RADIUS server status testing. You do not need to configure the test username and password to issue test messages to monitor RADIUS servers. You can use the default test username (test) and default password (test).

**Note**

We recommend that the test username not be the same as an existing username in the RADIUS database for security reasons.

About Validating a RADIUS Server

As of Cisco SAN-OS Release 3.0(1), you can periodically validate a RADIUS server. The switch sends a test authentication to the server using the username and password that you configure. If the server does not respond to the test authentication, then the server is considered non responding.

**Note**

For security reasons we recommend that you do not use a username that is configured on your RADIUS server as a test username.

You can configure this option to test the server periodically, or you can run a one-time only test.

About Vendor-Specific Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific attributes (VSAs) between the network access server and the RADIUS server. The IETF uses attribute 26. VSAs allow vendors to support their own extended attributes that are not suitable for general use. The Cisco RADIUS implementation supports one vendor-specific option using the format recommended in the specification. The Cisco vendor ID is 9, and the supported option is vendor type 1, which is named **cisco-avpair**. The value is a string with the following format:

```
protocol : attribute separator value *
```

Where **protocol** is a Cisco attribute for a particular type of authorization, **separator** is = (equal sign) for mandatory attributes, and * (asterisk) is for optional attributes.

When you use RADIUS servers to authenticate yourself to a Cisco MDS 9000 Family switch, the RADIUS protocol directs the RADIUS server to return user attributes, such as authorization information, along with authentication results. This authorization information is specified through VSAs.

VSA Format

The following VSA protocol options are supported by the Cisco NX-OS software:

- **Shell** protocol—Used in Access-Accept packets to provide user profile information.
- **Accounting** protocol—Used in Accounting-Request packets. If a value contains any white spaces, it should be put within double quotation marks.

The following attributes are supported by the Cisco NX-OS software:

- **roles**—This attribute lists all the roles to which the user belongs. The value field is a string storing the list of group names delimited by white space. For example, if you belong to roles **vsan-admin** and **storage-admin**, the value field would be **“vsan-admin storage-admin”**. This subattribute is sent in the VSA portion of the Access-Accept frames from the RADIUS server, and it can only be used with the shell protocol value. These are two examples using the roles attribute:

```
shell:roles="network-admin vsan-admin"
```

```
shell:roles*"network-admin vsan-admin"
```

When an VSA is specified as **shell:roles*"network-admin vsan-admin"**, this VSA is flagged as an optional attribute, and other Cisco devices ignore this attribute.

- **accountinginfo**—This attribute stores additional accounting information besides the attributes covered by a standard RADIUS accounting protocol. This attribute is only sent in the VSA portion of the Account-Request frames from the RADIUS client on the switch, and it can only be used with the accounting protocol-related PDUs.

Specifying SNMPv3 on AAA Servers

The vendor/custom attribute **cisco-av-pair** can be used to specify user's role mapping using the format:

```
shell:roles="roleA roleB ..."
```

If the role option in the **cisco-av-pair** attribute is not set, the default user role is network-operator.

The VSA format optionally specifies your SNMPv3 authentication and privacy protocol attributes also as follows:

```
shell:roles="roleA roleB..." snmpv3:auth=SHA priv=AES-128
```

The SNMPv3 authentication protocol options are SHA and MD5. The privacy protocol options are AES-128 and DES. If these options are not specified in the **cisco-av-pair** attribute on the ACS server, MD5 and DES are used by default.

One-Time Password Support

A one-time password (OTP) is a password that is valid for a single login session or transaction. OTPs avoid a number of shortcomings that are associated with usual (static) passwords. The most vital shortcoming that is addressed by OTPs is that, in contrast to static passwords, they are not at risk to replay attacks. If an intruder manages to record an OTP that was already used to log into a service or to conduct an operation, it will not be misused because it is no longer valid.

One-time password applies only to RADIUS and TACACS protocol daemons. In the case of the RADIUS protocol daemon, there is no configuration required from the switch side. In the case of the TACACS protocol, ASCII authentication mode needs to be enabled using the following command.

```
aaa authentication login ascii-authentication
```

About TACACS+

TACACS+ is a client/server protocol that uses TCP (TCP port 49) for transport requirements. All switches in the Cisco MDS 9000 Family provide centralized authentication using the TACACS+ protocol. The TACACS+ has the following advantages over RADIUS authentication:

- Provides independent, modular AAA facilities. Authorization can be done without authentication.
- Uses the TCP transport protocol to send data between the AAA client and server, making reliable transfers with a connection-oriented protocol.
- Encrypts the entire protocol payload between the switch and the AAA server to ensure higher data confidentiality. The RADIUS protocol only encrypts passwords.

About TACACS+ Server Default Configuration

DCNM-SAN allows you to set up a default configuration that can be used for any TACACS+ server that you configure the switch to communicate with. The default configuration includes:

- Encryption type
- Preshared key
- Timeout value
- Number of retransmission attempts
- Allowing the user to specify a TACACS+ server at login

About the Default TACACS+ Server Encryption Type and Preshared Key

You need to configure the TACACS+ preshared key to authenticate the switch to the TACACS+ server. The length of the key is restricted to 64 characters and can include any printable ASCII characters (white spaces are not allowed). You can configure a global key to be used for all TACACS+ server configurations on the switch.

You can override this global key assignment by explicitly using the **key** option when configuring an individual TACACS+ server.

About TACACS+ Servers

By default, the TACACS+ feature is disabled in all switches in the Cisco MDS 9000 Family. DCNM-SAN or Device Manager enables the TACACS+ feature automatically when you configure a TACACS+ server.

If a secret key is not configured for a configured server, a warning message is issued if a global key is not configured. If a server key is not configured, the global key (if configured) is used for that server.



Note Prior to Cisco MDS SAN-OS Release 2.1(2), you can use the dollar sign (\$) in the key but the key must be enclosed in double quotes, for example “k\$”. The percent sign (%) is not allowed. In Cisco MDS SAN-OS Release 2.1(2) and later, you can use the dollar sign (\$) without double quotes and the percent sign (%) in global secret keys.

You can configure global values for the secret key for all TACACS+ servers.



Note If secret keys are configured for individual servers, those keys override the globally configured key.

Password Aging Notification through TACACS+ Server

Password aging notification is initiated when the user authenticates to a Cisco MDS 9000 switch via a TACACS+ account. The user is notified when a password is about to expire or has expired. If the password has expired, user is prompted to change the password.



Note As of Cisco MDS SAN-OS Release 3.2(1), only TACACS+ supports password aging notification. If you try to use RADIUS servers by enabling this feature, RADIUS generates a SYSLOG message and authentication falls back to the local database.

Password aging notification facilitates the following:

- Password change—You can change your password by entering a blank password.
- Password aging notification—Notifies password aging. Notification happens only if the AAA server is configured and MSCHAP and MSCHAPv2 is disabled.
- Password change after expiration—Initiates password change after the old password expires. Initiation happens from the AAA server.



Note Password aging notification fails if you do not disable MSCHAP and MSCHAPv2 authentication.

To enable the password aging option in the AAA server, enter the following command:

```
aaa authentication login ascii-authentication
```

To determine whether or not password aging notification is enabled or disabled in the AAA server, enter the following command:

```
show aaa authentication login ascii-authentication
```

About Validating a TACACS+ Server

As of Cisco SAN-OS Release 3.0(1), you can periodically validate a TACACS+ server. The switch sends a test authentication to the server using the test username and test password that you configure. If the server does not respond to the test authentication, then the server is considered nonresponding.



Note We recommend that you do not configure the test user on your TACACS+ server for security reasons.

You can configure this option to test the server periodically, or you can run a one-time only test.

Periodically Validating a TACACS+ Server

To configure the switch to periodically test a TACACS+ server using DCNM-SAN, see the [Configuring the RADIUS, TACACS+, and LDAP Server, on page 635](#).

About Users Specifying a TACACS+ Server at Login

By default, an MDS switch forwards an authentication request to the first server in the TACACS+ server group. You can configure the switch to allow the user to specify which TACACS+ server to send the authenticate request. If you enable this feature, the user can log in as *username@hostname*, where the *hostname* is the name of a configured TACACS+ server.

Supported TACACS+ Server Parameters

The Cisco NX-OS software currently supports the following parameters for the listed TACACS+ servers:

- TACACS+

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

- Cisco ACS TACACS+

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

- Open TACACS+

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

About Bypassing a Nonresponsive Server

As of Cisco SAN-OS Release 3.0(1), you can bypass a nonresponsive AAA server within a server group. If the switch detects a nonresponsive server, it will bypass that server when authenticating users. Use this feature to minimize login delays caused by a faulty server. Instead of sending a request to a nonresponsive server and waiting for the authentication request to timeout, the switch sends the authentication request to the next server in the server group. If there are no other responding servers in the server group, the switch continues to attempt authentications against the nonresponsive server.

AAA Server Distribution

Configuration for RADIUS and TACACS+ AAA on an MDS switch can be distributed using the Cisco Fabric Services (CFS). The distribution is disabled by default (see the Cisco MDS 9000 Family NX-OS System Management Configuration Guide).

After enabling the distribution, the first server or global configuration starts an implicit session. All server configuration commands entered thereafter are stored in a temporary database and applied to all switches in the fabric (including the originating one) when you explicitly commit the database. The various server and global parameters are distributed, except the server and global keys. These keys are unique secrets to a switch and should not be shared with other switches.



Note Server group configurations are not distributed.



Note For an MDS switch to participate in AAA server configuration distribution, it must be running Cisco MDS SAN-OS Release 2.0(1b) or later, or Cisco NX-OS Release 4.1(1).

Starting a Distribution Session on a Switch

A distribution session starts the moment you begin a RADIUS or TACACS+ server or global configuration. For example, the following tasks start an implicit session:

- Specifying the global timeout for RADIUS servers.
- Specifying the global timeout for TACACS+ servers.



Note After you issue the first configuration command related to AAA servers, all server and global configurations that are created (including the configuration that caused the distribution session start) are stored in a temporary buffer, not in the running configuration.

CHAP Authentication

Challenge Handshake Authentication Protocol (CHAP) is a challenge-response authentication protocol that uses the industry-standard Message Digest 5 (MD5) hashing scheme to encrypt the response. CHAP is used by various vendors of network access servers and clients. A server running routing and remote access supports CHAP so that remote access clients that require CHAP are authenticated. CHAP is supported as an authentication method in this release.

MSCHAP Authentication

Microsoft Challenge Handshake Authentication Protocol (MSCHAP) is the Microsoft version of CHAP.

Cisco MDS 9000 Family switches allow user logins to perform remote authentication using different versions of MSCHAP. MSCHAP is used for authentication on a RADIUS or TACACS+ server, while MSCHAPv2 is used for authentication on a RADIUS server.

About Enabling MSCHAP

By default, the switch uses Password Authentication Protocol (PAP) authentication between the switch and the remote server. If you enable MSCHAP, you need to configure your RADIUS server to recognize the MSCHAP vendor-specific attributes. See the [About Vendor-Specific Attributes, on page 626](#). [Table 76: MSCHAP RADIUS Vendor-Specific Attributes, on page 631](#) shows the RADIUS vendor-specific attributes required for MSCHAP.

Table 76: MSCHAP RADIUS Vendor-Specific Attributes

Vendor-ID Number	Vendor-Type Number	Vendor-Specific Attribute	Description
311	11	MSCHAP-Challenge	Contains the challenge sent by an AAA server to an MSCHAP user. It can be used in both Access-Request and Access-Challenge packets.
211	11	MSCHAP-Response	Contains the response value provided by an MS-CHAP user in response to the challenge. It is only used in Access-Request packets.

Local AAA Services

The system maintains the username and password locally and stores the password information in encrypted form. You are authenticated based on the locally stored user information.

Accounting Services

Accounting refers to the log information that is kept for each management session in a switch. This information may be used to generate reports for troubleshooting and auditing purposes. Accounting can be implemented locally or remotely (using RADIUS). The default maximum size of the accounting log is 250,000 bytes and cannot be changed.



Tip The Cisco MDS 9000 Family switch uses interim-update RADIUS accounting-request packets to communicate accounting log information to the RADIUS server. The RADIUS server must be appropriately configured to log the information communicated in these packets. Several servers typically have log update/watchdog packets flags in the AAA client configuration. Turn on this flag to ensure proper RADIUS accounting.



Note Configuration operations are automatically recorded in the accounting log if they are performed in configuration mode. Additionally, important system events (for example, configuration save and system switchover) are also recorded in the accounting log.

Defining Roles on the Cisco Secure ACS 5.x GUI

Enter the following in the GUI under Policy Elements:

Table 77: Role Definitions

Attribute	Requirement	Value
shell:roles	Optional	network-admin

Defining Custom Attributes for Roles

Cisco MDS 9000 Family switches use the TACACS+ custom attribute for service shells to configure roles to which a user belongs. TACACS+ attributes are specified in **name=value** format. The attribute name for this custom attribute is **cisco-av-pair**. The following example illustrates how to specify roles using this attribute:

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

You can also configure optional custom attributes to avoid conflicts with non-MDS Cisco switches using the same AAA servers.

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

Additional custom attribute shell:roles are also supported:

```
shell:roles="network-admin vsan-admin
"
```

or

```
shell:roles*"network-admin vsan-admin"
```



Note

TACACS+ custom attributes can be defined on an Access Control Server (ACS) for various services (for example, shell). Cisco MDS 9000 Family switches require the TACACS+ custom attribute for the service shell to be used for defining roles.

Guidelines and Limitations

This section has the following topics:

Remote Authentication Guidelines

If you prefer using remote AAA servers, follow these guidelines:

- A minimum of one AAA server should be IP reachable.
- Be sure to configure a desired local AAA policy as this policy is used if all AAA servers are not reachable.
- AAA servers are easily reachable if an overlay Ethernet LAN is attached to the switch (see the IP Services Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Configuration Guide). We recommend this method.

SAN networks connected to the switch should have at least one gateway switch connected to the Ethernet LAN reaching the AAA servers.

Guidelines and Limitations for LDAP

LDAP has the following guidelines and limitations:

- You can configure a maximum of 64 LDAP servers on the Cisco NX-OS device.
- Cisco NX-OS supports only LDAP version 3.
- Cisco NX-OS supports only these LDAP servers:
 - OpenLDAP
 - Microsoft Active Directory
- LDAP over Secure Sockets Layer (SSL) supports only SSL version 3 and Transport Layer Security (TLS) version 1.
- If you have a user account configured on the local Cisco NX-OS device that has the same name as a remote user account on an AAA server, the Cisco NX-OS software applies the user roles for the local user account to the remote user, not the user roles configured on the AAA server.

Merge Guidelines for RADIUS and TACACS+ Configurations

The RADIUS and TACACS+ server and global configuration are merged when two fabrics merge. The merged configuration is applied to CFS distribution-enabled switches.

When merging the fabric, be aware of the following conditions:

- The server groups are not merged.
- The server and global keys are not changed during the merge.
- The merged configuration contains all servers found on all CFS enabled switches.
- The timeout and retransmit parameters of the merged configuration are the largest values found per server and global configuration.



Note

Test parameter will be distributed via CFS for TACACS+ Daemon only. If the fabric contains only Cisco NX-OS Release 5.0 devices, then the test parameters will be distributed. If the fabric contains devices running Release 5.0 and some running Release 4.x, the test parameters are not distributed.



Caution

If there is a conflict between two switches in the server ports configured, the merge fails.

Default Settings

[Table 78: Default LDAP Parameter Settings](#), on page 633 lists the default settings for LDAP parameters.

Table 78: Default LDAP Parameter Settings

Parameters	Default
LDAP	Disabled
LDAP authentication method	First search and then bind
LDAP authentication mechanism	Plain

Parameters	Default
Dead-interval time	0 minutes
Timeout interval	5 seconds
Idle timer interval	60 minutes
Periodic server monitoring username	test
Periodic server monitoring password	Cisco

[Table 79: Default Switch Security Settings](#) , on page 634 lists the default settings for all switch security features in any switch.

Table 79: Default Switch Security Settings

Parameters	Default
Roles in Cisco MDS switches	Network operator (network-operator)
AAA configuration services	Local
Authentication port	1812
Accounting port	1813
Preshared key communication	Clear text
RADIUS server timeout	1 (one) second
RADIUS server retries	Once
Authorization	Disabled
aaa user default role	enabled
RADIUS server directed requests	Disabled
TACACS+	Disabled
TACACS+ servers	None configured
TACACS+ server timeout	5 seconds
TACACS+ server directed requests	Disabled
AAA server distribution	Disabled
Accounting log size	250 KB

Configuring the RADIUS, TACACS+, and LDAP Server

Cisco MDS 9000 Family switches can use the RADIUS protocol to communicate with remote AAA servers. You can configure multiple RADIUS servers and server groups and set timeout and retry counts.

RADIUS is a distributed client/server protocol that secures networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco MDS 9000 Family switches and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.

This section defines the RADIUS operation, identifies its network environments, and describes its configuration possibilities.

A Cisco MDS switch uses the Terminal Access Controller Access Control System Plus (TACACS+) protocol to communicate with remote AAA servers. You can configure multiple TACACS+ servers and set timeout values.

This section includes the following topics:

Authorizing and Authenticating the Switch

To authorize and authenticate the switch, follow these steps:

Procedure

-
- | | |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Log in to the required switch in the Cisco MDS 9000 Family, using the Telnet, SSH, DCNM-SAN or Device Manager, or console login options. |
| Step 2 | <p>When you have configured server groups using the server group authentication method, an authentication request is sent to the first AAA server in the group.</p> <ul style="list-style-type: none">• If the AAA server fails to respond, then the next AAA server is contacted and so on until the remote server responds to the authentication request.• If all AAA servers in the server group fail to respond, then the servers in the next server group are contacted.• If all configured methods fail, then by default local database is used for authentication. The next section will describe the way to disable this fallback. |
| Step 3 | <p>When you are successfully authenticated through a remote AAA server, then the following possible actions are taken:</p> <ul style="list-style-type: none">• If the AAA server protocol is RADIUS, then user roles specified in the cisco-av-pair attribute are downloaded with an authentication response.• If the AAA server protocol is TACACS+, then another request is sent to the same server to get the user roles specified as custom attributes for the shell.• If user roles are not successfully retrieved from the remote AAA server, then the user is assigned the network-operator role if the show aaa user default-role command is enabled. You are denied access if this command is disabled. |

- Step 4** When your user name and password are successfully authenticated locally, you are allowed to log in, and you are assigned the roles configured in the local database.

Configuring Fallback Mechanism for Authentication

You can enable or disable fallback to the local database in case the remote authentication is set and all of the AAA servers are unreachable (authentication error). The fallback is set to local by default in case of an authentication error. You can disable this fallback for both console and SSH or Telnet login. Disabling this fallback tightens the authentication security.

To configure the fallback mechanism, follow this step:

Procedure

Enter the `show run aaa all` command to verify that the default fallback is enabled for both the default and console login.

Disabling fallback will print a warning message.

Caution If fallback is disabled for both the default and console, remote authentication is enabled and servers are unreachable and then the switch will be locked.

Configuring AAA Server Monitoring Parameters Globally

The AAA server monitoring parameters can be configured globally for all servers or individually for a specific server. This section explains how the global configuration can be set. The global configurations will apply to all servers that do not have individual monitoring parameters defined. For any server, the individual test parameter defined for that particular server will always get precedence over the global settings.

Procedure

	Command or Action	Purpose
Step 1	Configure the global monitoring parameters for RADIUS servers	

Configuring the Default RADIUS Server Encryption Type and Preshared Key

To configure the default RADIUS server encryption type and preshared key, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	Expand Switches > Security > AAA , and then select RADIUS .	You see the RADIUS configuration in the Information pane.
Step 2	Click the Defaults tab.	

	Command or Action	Purpose
Step 3	Click the Apply Changes icon to save the changes.	
Step 4	Set the key in the Auth Key field.	
Step 5	Select plain or encrypted from the AuthType drop-down menu.	

Setting the Default RADIUS Server Timeout Interval and Retransmits

By default, a switch retries transmission to a RADIUS server only once before reverting to local authentication. You can increase this number up to a maximum of five retries per server. You can also configure the timeout value for the RADIUS server.

To configure the number of retransmissions and the time between retransmissions to the RADIUS servers, follow these steps:

Procedure

-
- Step 1** Expand **Switches > Security > AAA**, and then select **RADIUS**.
You see the RADIUS configuration in the Information pane.
 - Step 2** Choose the **Defaults** tab.
You see the RADIUS default settings.
 - Step 3** Fill in the Timeout and Retransmits fields for authentication attempts.
 - Step 4** Click the **Apply Changes** icon to save the changes.
-

Configuring an LDAP Server

To configure an LDAP server and all of its options, follow these steps:

Procedure

-
- Step 1** Expand **Switches > Security > AAA**, and then select **LDAP**.
You see the LDAP configuration in the Information pane.
 - Step 2** Click the **Servers** tab.
You see any existing RADIUS servers.
 - Step 3** Click **Create Row** to add a new LDAP server.
You see the Create LDAP Server dialog box.

VDC Setup Wizard

Step 6 of 6: Management of VDC

Specify the parameters to enable management of VDC.

IPv4 Management Interface

IPv4 Address:

☒ Prefix Length:

☐ Netmask:

Default Gateway:

IPv6 Management Interface

IPv6 Address:

Prefix Length:

Default Gateway:

SSH

☒ Enable SSH Server

SSH Key Type:

SSH Key Length: (1024-2048)

☒ Discover the VDC

At the end of VDC creation, DCNM will start discovering the VDC with the below mentioned Credentials.

User Name:

Password:

< Back Next > Finish Cancel

- Step 4** Select the switches that you want to assign as LDAP servers.
- Step 5** Assign an index number to identify the LDAP server.
- Step 6** Select the IP address type for the LDAP server.
- Step 7** Fill in the IP address or name for the LDAP server.
- Step 8** (Optional) Modify the authentication and accounting ports used by this LDAP server.
- Step 9** Select the appropriate key type for the LDAP server.
- Step 10** Select the TimeOut value in seconds. The valid range is 0 to 60 seconds.
- Step 11** Select the number of times the switch tries to connect to an LDAP server(s) before reverting to local authentication.
- Step 12** Enter the test idle time interval value in minutes. The valid range is 1 to 1440 minutes.
- Step 13** Enter the test user with the default password. The default username is test.
- Step 14** Click **Create** to save these changes.

Creating LDAP Search Map

To create an LDAP search map, follow these steps:

Procedure

- Step 1** Expand **Switches > Security > AAA**, and then select **LDAP**.
You see the LDAP configuration in the Information pane.
- Step 2** Click the **Search Map** tab.

- Step 3** Click **Create Row** to add a new LDAP search map.
- Step 4** Enter the LDAP search map name for the Name field.
- Step 5** Select the appropriate search type for the Type field.
- Step 6** Enter the base domain name for the BaseDN field.
- Step 7** Enter the filter value for the Filter field.
- Step 8** Enter the attribute value for the Attribute field.
- Step 9** Click Create to save the changes.

Configuring a RADIUS Server

To configure a RADIUS server and all its options, follow these steps:

Procedure

- Step 1** Expand **Switches > Security > AAA**, and then select **RADIUS**.
You see the RADIUS configuration in the Information pane.
- Step 2** Click the **Servers** tab.
You see any existing RADIUS servers.
- Step 3** Click **Create Row** to add a new RADIUS server.
You see the Create RADIUS Server dialog box shown in the following figure.

Figure 84: Create RADIUS Server

- Step 4** Select the switches that you want to assign as RADIUS servers.
 - Step 5** Assign an index number to identify the RADIUS server.
 - Step 6** Select the IP address type for the RADIUS server.
 - Step 7** Fill in the IP address or name for the RADIUS server.
 - Step 8** (Optional) Modify the authentication and accounting ports used by this RADIUS server.
 - Step 9** Select the appropriate key type for the RADIUS server.
 - Step 10** Select the TimeOut value in seconds. The valid range is 0 to 60 seconds.
 - Step 11** Select the number of times the switch tries to connect to a RADIUS server(s) before reverting to local authentication.
 - Step 12** Enter the test idle time interval value in minutes. The valid range is 1 to 1440 minutes.
 - Step 13** Enter the test user with the default password. The default username is test.
 - Step 14** Click **Create** to save these changes.
-

Validating a RADIUS Server

To configure the switch to periodically test a RADIUS server, follow these steps:

Procedure

- Step 1** Expand **Switches > Security > AAA**, and then select **RADIUS**.
You see the RADIUS configuration in the Information pane.
 - Step 2** Click the **Servers** tab.
You see any existing RADIUS servers.
 - Step 3** Click **Create Row** to add a new RADIUS server.
You see the Create RADIUS Server dialog box.
 - Step 4** Fill in the IP address.
 - Step 5** Modify the authentication and accounting ports used by this RADIUS server.
 - Step 6** Fill in the TestUser field and, optionally, the TestPassword field. The default password for the test is **Cisco**.
 - Step 7** Set the IdleTime field for the time that the server is idle before you send a test authentication.
 - Step 8** Click **Create** to save these changes.
-

Allowing Users to Specify a RADIUS Server at Login

By default, an MDS switch forwards an authentication request to the first server in the RADIUS server group. You can configure the switch to allow the user to specify which RADIUS server to send the authenticate request by enabling the directed request option. If you enable this option, the user can log in as *username@hostname*, where the *hostname* is the name of a configured RADIUS server.

To allow users logging into an MDS switch to select a RADIUS server for authentication, follow these steps:

To allow users logging into an MDS switch to select a RADIUS server for authentication, follow these steps:

Procedure

- Step 1** Expand **Switches** > **Security** > **AAA**, and then select **RADIUS**.
You see the RADIUS configuration in the Information pane.
- Step 2** Click the **Defaults** tab.
You see the RADIUS default settings.
- Step 3** Check the **DirectedReq** check box for the RADIUS server.
- Step 4** Click the **Apply Changes** icon to save the changes.
-

Setting the Default TACACS+ Server Encryption Type and Preshared Key

To configure the default TACACS+ server encryption type and preshared key, follow these steps:

Procedure

- Step 1** Expand **Switches** > **Security** > **AAA**, and then select **TACACS+**.
You see the TACACS+ configuration in the Information pane.
- Step 2** If the Defaults tab is dimmed, click the **CFS** tab.
- Step 3** Click the **Defaults** tab.
You see the TACACS+ default settings.
- Step 4** Select **plain** or **encrypted** from the AuthType drop-down menu and set the key in the Auth Key field.
- Step 5** Click the **Apply Changes** icon to save the changes.
-

Setting the Default TACACS+ Server Timeout Interval and Retransmits

By default, a switch retries a TACACS+ server only once. This number can be configured. The maximum is five retries per server. You can also configure the timeout value for the TACACS+ server.

To configure the number of retransmissions and the time between retransmissions to the TACACS+ servers, follow these steps:

Procedure

- Step 1** Expand **Switches** > **Security** > **AAA**, and then select **TACACS+**.
You see the TACACS+ configuration in the Information pane.
- Step 2** Click the **Defaults** tab. (If the **Defaults** tab is disabled, click the **CFS** tab first).
You see the TACACS+ default settings.

- Step 3** Supply values for the Timeout and Retransmits fields for authentication attempts.
- Step 4** Click the **Apply Changes** icon to save the changes.

Configuring a TACACS+ Server

To configure a TACACS+ server and all its options using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand **Switches > Security > AAA**, and then select **TACACS+**.
You see the TACACS+ configuration in the Information pane.
- Step 2** Click the **Servers** tab.
You see any existing TACACS+ servers.
- Step 3** Click **Create Row** to add a new TACACS+ server.
You see the Create TACACS+ Server dialog box as shown in the following figure.

Figure 85: Create TACACS+ Server Dialog Box

- Step 4** Select the switches that you want to assign as TACACS servers.
- Step 5** Assign an index number to identify the TACACS server.
- Step 6** Select the IP address type for the TACACS server.
- Step 7** Fill in the IP address or name for the TACACS server.

- Step 8** Modify the authentication and accounting ports used by this TACACS server.
 - Step 9** Select the appropriate key type for the TACACS server.
 - Step 10** Select the TimeOut value in seconds. The valid range is 0 to 60 seconds.
 - Step 11** Select the number of times the switch tries to connect to a TACACS server(s) before reverting to local authentication.
 - Step 12** Enter the test idle time interval value in minutes. The valid range is 1 to 1440 minutes.
 - Step 13** Enter the test user with the default password. The default username is test.
 - Step 14** Click **Create** to save these changes.
-

Allowing Users to Specify a TACACS+ Server at Login

To allow users logging into an MDS switch to select a TACACS+ server for authentication, follow these steps:

To configure the switch to allow users to specify a TACACS+ server at login using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand **Switches > Security > AAA**, and then select **TACACS+**.
You see the TACACS+ configuration in the Information pane.
 - Step 2** Click the **Defaults** tab.
You see the TACACS+ default settings.
 - Step 3** Check the **DirectedReq** check box.
 - Step 4** Click the **Apply Changes** icon to save the changes.
-

Clearing TACACS+ Server Statistics

You can clear all the TACACS+ server statistics using the clear tacacs-server statistics 10.1.2.3 command.

Configuring Server Groups

To configure a RADIUS, TACACS+, or LDAP server group using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand **Switches > Security**, and then select **AAA**.
You see the AAA configuration in the Information pane. If you do not see the screen, click the **Server Groups** tab.
You see the RADIUS, TACACS+, or LDAP server groups configured.
- Step 2** Click **Create Row** to create a server group.

You see the Create Server dialog box.

- Step 3** Click the **radius** radio button to add a RADIUS server group, the **tacacs+ radio button** to add a TACACS+ server group, and the **ldap** radio button to add a LDAP server group.
- Step 4** Supply server names for the ServerIdList field.
- Step 5** When you chose LDAP, enter the LDAP search map name for the LDAPSearchMapName.
- LDAPSSLMODE—Specifies if the TLS tunnel should be setup before binding with the LDAP server.
 - LDAPBindFirst—Specifies if the user bind should be completed before the search.
- Step 6** Click the plain radio button to select the plain authentication method, click the kerberos button to select the kerberos authentication method, and click md5digest to select the md5digest authentication method.
- Step 7** Enter the password for the LDAPComparePasswd field:
- LDAPCertDNBind—Specifies if the User Certification Bind needs to be checked while doing PKI SSH certificate authorization.
 - LDAPUserServerBind—Specifies if the User Server Bind should be checked as part of SSH PKI authorization.
- Step 8** Set the DeadTime field for the number of minutes that a server can be nonresponsive before it is marked as bypassed. See the [About Bypassing a Nonresponsive Server, on page 629](#).
- Step 9** Click **Create** to create this server group.
- The LDAP Server Group displays LDAP-specific parameters.
- Step 10** Click the **Applications** tab to assign this server group to an application.
- You can associate a server group with all applications or you can specify specific applications.
- Step 11** Click the General tab to assign the type of authentication to this server group.
- Check either the MSCHAP or MSCHAPv2 check box based on the type of server group.
- Step 12** Click the **Apply Changes** icon to save the changes.
- Once the LDAP Server group is created, the configuration information is displayed in two tabs:
- Server Groups—Displays common data shared by all AAA protocols (RADIUS, TACACS+, and LDAP).
 - LDAP Server Group—Displays only LDAP-specific protocols.

Enabling Radius Server Distribution

To enable RADIUS server distribution, follow these steps:

Procedure

- Step 1** Expand **Switches > Security > AAA**, and then select **RADIUS**.
- You see the RADIUS configuration in the Information pane.

- Step 2** Click the **CFS** tab. You see the RADIUS CFS configuration.
 - Step 3** Choose **enable** from the Admin drop-down list for all switches that you want to enable CFS for RADIUS.
 - Step 4** Click **Apply Changes** to distribute these changes through the fabric.
-

Enabling TACACS+ Server Distribution

To enable TACACS+ server distribution, follow these steps:

Procedure

- Step 1** Expand **Switches > Security > AAA**, and then select **TACACS+**.
You see the TACACS+ configuration in the Information pane.
 - Step 2** Click the **CFS** tab.
You see the TACACS+ CFS configuration.
 - Step 3** Choose **enable** from the Admin drop-down list for all switches that you want to enable CFS on for TACACS+.
 - Step 4** Click **Apply Changes** to distribute these changes through the fabric.
-

Committing the Distribution

To distribute a RADIUS or TACACS+ configuration, follow these steps:

Procedure

- Step 1** Expand **Switches > Security > AAA**, and then select either **RADIUS** or **TACACS+**. You see the RADIUS or TACACS+ configuration in the Information pane.
 - Step 2** Click the **CFS** tab. You see the RADIUS or TACACS+ CFS configuration.
 - Step 3** Choose **commitChanges** in the Config Action drop-down list for all switches that you want to enable CFS for RADIUS or TACACS+.
 - Step 4** Click **Apply Changes** to distribute the changes through the fabric.
-

Discarding the Distribution Session

To discard RADIUS or TACACS+ distribution, follow these steps:

Procedure

- Step 1** Expand **Switches > Security > AAA**, and then select either **RADIUS** or **TACACS+**. You see either the RADIUS or TACACS+ configuration in the Information pane.
- Step 2** Click the **CFS** tab. You see either the RADIUS or TACACS+ CFS configuration.

- Step 3** Choose **abort** from the Config Action drop-down list for each switch that should discard the pending RADIUS or TACACS+ distribution.
- Step 4** Click **Apply Changes**.

What to do next

Clearing Sessions

To clear the ongoing CFS distribution session (if any) and to unlock the fabric for the RADIUS feature, enter the **clear radius session** command from any switch in the fabric.

```
switch# clear radius session
```

To clear the ongoing CFS distribution session (if any) and to unlock the fabric for the TACACS+ feature, enter the **clear tacacs+ session** command from any switch in the fabric.

```
switch# clear tacacs+ session
```

To clear a RADIUS or TACACS+ distribution, follow these steps:

Procedure

- Step 1** Expand **Switches > Security > AAA** and then select either **RADIUS** or **TACACS+**.
You see either the RADIUS or TACACS+ configuration in the Information pane.
- Step 2** Choose the **CFS** tab. You see either the RADIUS or TACACS+ CFS configuration.
- Step 3** Choose **clear** from the Config Action drop-down list for each switch that should clear the pending RADIUS or TACACS+ distribution.
- Step 4** Click **Apply Changes**.

Enabling MSCHAP Authentication



Note Password aging, MSCHAPv2, and MSCHAP authentication can fail if one of these authentication is not disabled.



Note A warning message is issued when you execute a command to enable MSCHAPv2 authentication on the TACACS+ server, and the configuration fails.

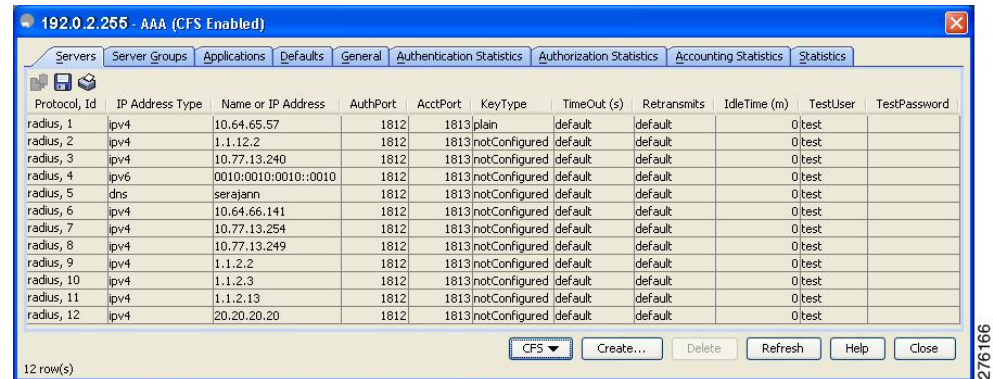
To enable MSCHAP authentication using Device Manager, follow these steps:

Procedure

Step 1 Click **Security > AAA**.

You see the AAA configuration in the Information pane.


Figure 86: AAA Configuration in Device Manager



Step 2 Click the **General** tab.

You see the MSCHAP configuration.

Figure 87: MSCHAP Configuration



Step 3 Check the **AuthTypeMSCHAP** or **AuthTypeMSCHAPv2** check box to use MSCHAP or MSCHAPv2 to authenticate users on the switch.

Step 4 Click **Apply Changes** to save the changes.

Configuring Cisco Access Control Servers

The Cisco Access Control Server (ACS) uses TACACS+ and RADIUS protocols to provide AAA services that ensure a secure environment. When using the AAA server, user management is normally done using Cisco ACS. The following figures display ACS server user setup configurations for network-admin roles and multiple roles using either RADIUS or TACACS+.

Figure 88: Configuring the network-admin Role When Using RADIUS

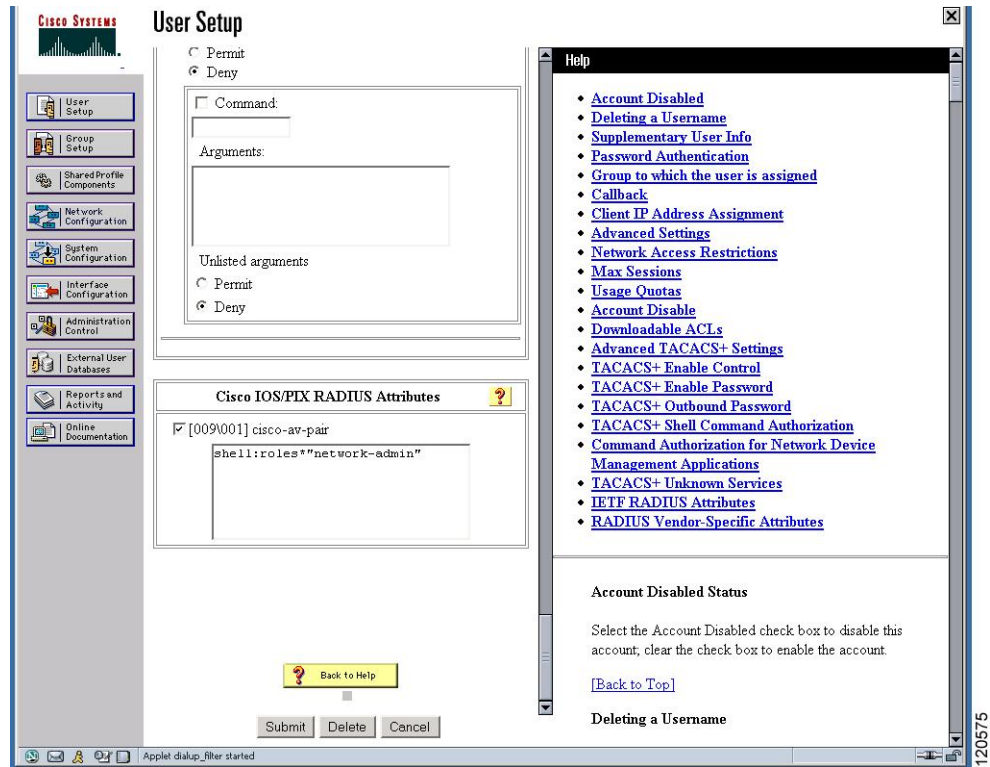


Figure 89: Configuring Multiple Roles with SNMPv3 Attributes When Using RADIUS

The screenshot shows the CiscoSecure ACS web interface for User Setup. The browser window title is "CiscoSecure ACS - Cisco Systems, Inc." and the address bar shows "http://10.76.100.108:2691/index2.htm". The interface includes a left sidebar with navigation links: User Setup, Group Setup, Shared Profile Components, Network Configuration, System Configuration, Interface Configuration, Administration Control, External User Databases, Reports and Activity, and Online Documentation. The main content area is titled "User Setup" and contains the following sections:

- Per User Command Authorization:** Includes "Unmatched Cisco IOS commands" with radio buttons for "Permit" and "Deny" (selected). Below is a "Command:" text field and an "Arguments:" text area. At the bottom, "Unlisted arguments" has radio buttons for "Permit" and "Deny" (selected).
- Cisco IOS/PIX RADIUS Attributes:** A checkbox labeled "[009/001] cisco-av-pair" is checked. Below it is a text area containing the configuration:


```
shell:roles="Role1 Role3 Role5
Role?"snmpv3:auth=MD5 priv=DES
```

At the bottom of the main content area are buttons for "Submit", "Delete", and "Cancel". On the right side, there is a "Help" panel with a list of links: Account Disabled, Deleting a Username, Supplementary User Info, Password Authentication, Group to which the user is assigned, Callback, Client IP Address Assignment, Advanced Settings, Network Access Restrictions, Max Sessions, Usage Quotas, Account Disable, Downloadable ACLs, Advanced TACACS+ Settings, TACACS+ Enable Control, TACACS+ Enable Password, TACACS+ Outbound Password, TACACS+ Shell Command Authorization, Command Authorization for Network Device Management Applications, TACACS+ Unknown Services, IETF RADIUS Attributes, and RADIUS Vendor-Specific Attributes. Below the links, the "Account Disabled Status" section states: "Select the Account Disabled check box to disable this account; clear the check box to enable the account." with a "[Back to Top]" link. The "Deleting a Username" section is partially visible at the bottom.

Figure 90: Configuring the network-admin Role with SNMPv3 Attributes When Using TACACS+

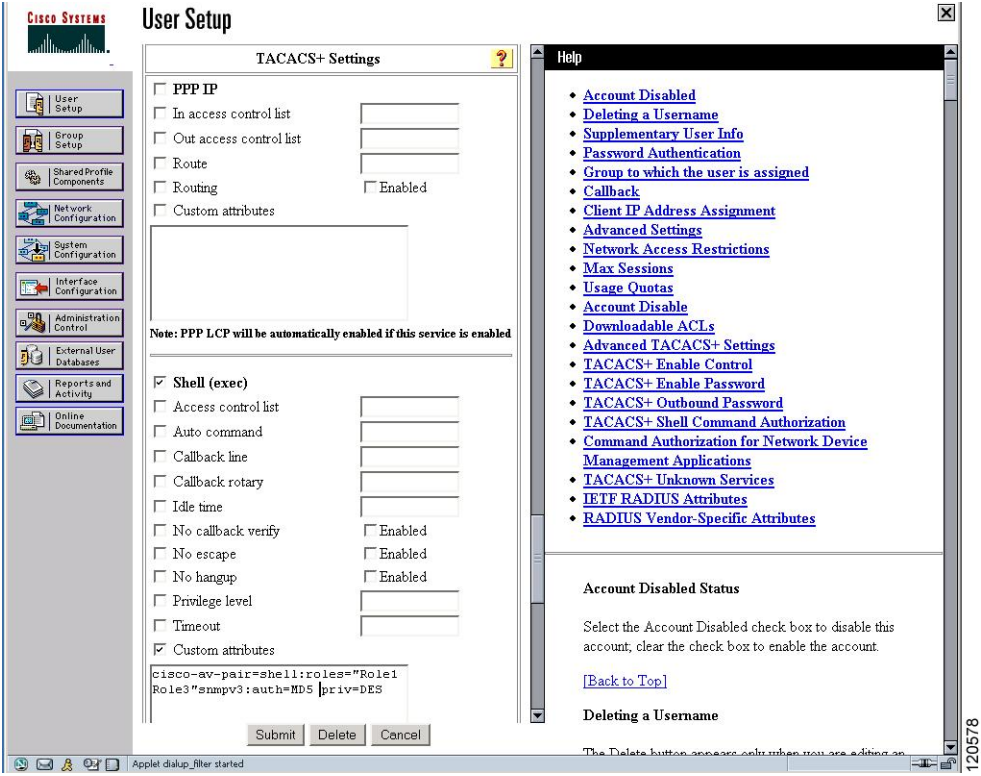


Figure 91: Configuring Multiple Roles with SNMPv3 Attributes When Using TACACS+

Verifying RADIUS and TACACS+ Configuration

To display the RADIUS and TACACS+ configuration information, perform one of the following tasks:

Command	Purpose
show show aaa authorization all	Displays aaa Authorization Information Details.
show aaa user default-role	Displays Default User Role for Remote Authentication
show radius-server directed-request	Display the RADIUS directed request configuration.
show tacacs-server directed-request	Display the TACACS+ directed request configuration.
show radius-server groups	Verify the configured server group order.
show aaa authentication login chap	Display the CHAP authentication configuration.
show aaa authentication login mschap	Display the MSCHAP authentication configuration.
show aaa authentication login mschapv2	Display the MSCHAPv2 authentication configuration.
show radius-server	Displays Configured RADIUS Information.
show radius-server groups	Displays Configured RADIUS Server-Group Order.

Command	Purpose
show radius-server statistics 10.1.3.2	Displays RADIUS Server Statistics.
show tacacs-server	Displays Configured TACACS+ Server Information.
show aaa authentication	Displays AAA Authentication Information.
show aaa authentication login error-enable	Displays AAA Authentication Login Information.
show tacacs-server groups	Displays Configured TACACS+ Server Groups.
show aaa groups	Displays All AAA Server Groups.
show tacacs-server statistics 10.1.2.3	Displays TACACS+ Server Statistics.
show radius distribution status	Displays the distribution status on the CFS tab.
show tacacs+ distribution status	Displays the session status once the implicit distribution session has started.
show radius pending-diff	Displays the RADIUS or TACACS+ global and/or server configuration stored in the temporary buffer.
show tacacs+ pending-diff	Displays the TACACS+ global and/or server configuration stored in the temporary buffer.
show radius distribution status	Displays the RADIUS Fabric Merge Status.
show tacacs+ distribution status	Displays the TACACS+ Fabric Merge Status.
show accounting log	Displays the Accounting Log Information.
show aaa authentication	Displays Authentication Information.
show accounting config	Displays Two Samples of Configured Accounting Parameters.
show accounting log 60000	Displays 60,000 Bytes of the Accounting Log.
show accounting log	Displays the Entire Log File.

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference*.

Displaying RADIUS Server Statistics

To display RADIUS server statistics, follow these steps:

Procedure

Step 1 Expand **Switches > Security > AAA**, and then select **RADIUS**.

You see the RADIUS configuration in the Information pane.

- Step 2** Click the **Statistics** tab.
You see the RADIUS server statistics.
-

Displaying TACACS+ Server Statistics

To display TACACS+ server statistics, follow these steps:

Procedure

- Step 1** Expand **Switches > Security > AAA**, and then select **TACACS+**.
You see the TACACS+ configuration in the Information pane.
- Step 2** Choose the **Statistics** tab.
You see the TACACS+ server statistics.
-

Displaying the Pending Configuration to be Distributed

To display the RADIUS or TACACS+ global and/or server configuration stored in the temporary buffer, follow these steps:

Procedure

- Step 1** Expand **Switches > Security > AAA**, and then select **RADIUS** or select **TACACS+**.
- Step 2** Click the CFS tab.
You see the distribution status on the CFS tab.
- Step 3** Click the **pending** or **running** radio button.
- Step 4** Click **Apply Changes** to save the changes.
- Step 5** Click the **Servers** tab to view the pending or running configuration.
-

What to do next

Use the **show radius distribution status** command to view the status of the RADIUS fabric merge.

Configuration Examples for LDAP

The following example shows how to configure an LDAP server host and server group:

```
feature ldap
ldap-server host 10.10.2.2 enable-ssl
aaa group server ldap LdapServer
```

```
server 10.10.2.2
exit
show ldap-server
show ldap-server groups
```

The following example shows how to configure an LDAP search map:

```
ldap search-map s0
userprofile attribute-name description search-filter
(&(objectClass=inetOrgPerson)(cn=$userid)) base-DN dc=acme,dc=com
exit
show ldap-search-map
```

The following example shows how to configure AAA authorization with certificate authentication for an LDAP server:

```
aaa authorization ssh-certificate default group LDAPServer1 LDAPServer2
exit
show aaa authorization
```




CHAPTER 30

Configuring Certificate Authorities and Digital Certificates

- [Configuring Certificate Authorities and Digital Certificates, on page 655](#)

Configuring Certificate Authorities and Digital Certificates

This chapter includes the following topics:

Information About Certificate Authorities and Digital Certificates

Public Key Infrastructure (PKI) support provides the means for the Cisco MDS 9000 Family switches to obtain and use digital certificates for secure communication in the network. PKI support provides manageability and scalability for IPsec/IKE and SSH.

Certificate Authorities (CAs) manage certificate requests and issue certificates to participating entities such as hosts, network devices, or users. The CAs provide centralized key management for the participating entities.

Digital signatures, based on public key cryptography, digitally authenticate devices and individual users. In public key cryptography, such as the RSA encryption system, each device or user has a key-pair containing both a private key and a public key. The private key is kept secret and is known only to the owning device or user only. However, the public key is known to everybody. The keys act as complements. Anything encrypted with one of the keys can be decrypted with the other. A signature is formed when data is encrypted with a sender's private key. The receiver verifies the signature by decrypting the message with the sender's public key. This process relies on the receiver having a copy of the sender's public key and knowing with a high degree of certainty that it really does belong to the sender and not to someone pretending to be the sender.

This section includes the following topics:

Purpose of CAs and Digital Certificates

CAs manage certificate requests and issue certificates to participating entities such as hosts, network devices, or users. The CAs provide centralized key management for the participating entities.

Digital signatures, based on public key cryptography, digitally authenticate devices and individual users. In public key cryptography, such as the RSA encryption system, each device or user has a key-pair containing both a private key and a public key. The private key is kept secret and is known only to the owning device or user only. However, the public key is known to everybody. The keys act as complements. Anything encrypted with one of the keys can be decrypted with the other. A signature is formed when data is encrypted with a

sender's private key. The receiver verifies the signature by decrypting the message with the sender's public key. This process relies on the receiver having a copy of the sender's public key and knowing with a high degree of certainty that it really does belong to the sender and not to someone pretending to be the sender.

Digital certificates link the digital signature to the sender. A digital certificate contains information to identify a user or device, such as the name, serial number, company, department, or IP address. It also contains a copy of the entity's public key. The certificate is itself signed by a CA, a third party that is explicitly trusted by the receiver to validate identities and to create digital certificates.

To validate the signature of the CA, the receiver must first know the CA's public key. Normally this process is handled out-of-band or through an operation done at installation. For instance, most web browsers are configured with the public keys of several CAs by default. The Internet Key Exchange (IKE), an essential component of IPsec, can use digital signatures to scalably authenticate peer devices before setting up security associations.

Trust Model, Trust Points, and Identity CAs

The trust model used in PKI support is hierarchical with multiple configurable trusted CAs. Each participating entity is configured with a list of CAs to be trusted so that the peer's certificate obtained during the security protocol exchanges can be verified, provided it has been issued by one of the locally trusted CAs. To accomplish this, the CA's self-signed root certificate (or certificate chain for a subordinate CA) is locally stored. The process of securely obtaining a trusted CA's root certificate (or the entire chain in the case of a subordinate CA) and storing it locally is called *CA authentication* and is a mandatory step in trusting a CA.

The information about a trusted CA that is locally configured is called the *trust point* and the CA itself is called a *trust point CA*. This information consists of CA certificate (or certificate chain in case of a subordinate CA) and the certificate revocation checking information.

The MDS switch can also enroll with a trust point to obtain an identity certificate (for example, for IPsec/IKE). This trust point is called an *identity CA*.

RSA Key-Pairs and Identity Certificates

You can generate one or more RSA key-pairs and associate each RSA key-pair with a trust point CA where the MDS switch intends to enroll to obtain an identity certificate. The MDS switch needs only one identity per CA, which consists of one key-pair and one identity certificate per CA.

Cisco MDS NX-OS allows you to generate RSA key-pairs with a configurable key size (or modulus). The default key size is 512. You can also configure an RSA key-pair label. The default key label is the switch fully qualified domain name (FQDN).

The following list summarizes the relationship between trust points, RSA key-pairs, and identity certificates:

- A trust point corresponds to a specific CA that the MDS switch trusts for peer certificate verification for any application (such as IKE or SSH).
- An MDS switch can have many trust points and all applications on the switch can trust a peer certificate issued by any of the trust point CAs.
- A trust point is not restricted to a specific application.
- An MDS switch enrolls with the CA corresponding to the trust point to obtain an identity certificate. You can enroll your switch with multiple trust points thereby obtaining a separate identity certificate from each trust point. The identity certificates are used by applications depending upon the purposes specified in the certificate by the issuing CA. The purpose of a certificate is stored in the certificate as certificate extensions.
- When enrolling with a trust point, you must specify an RSA key-pair to be certified. This key-pair must be generated and associated to the trust point before generating the enrollment request. The association

between the trust point, key-pair, and identity certificate is valid until it is explicitly removed by deleting the certificate, key-pair, or trust point.

- The subject name in the identity certificate is the fully qualified domain name for the MDS switch.
- You can generate one or more RSA key-pairs on a switch and each can be associated to one or more trust points. But no more than one key-pair can be associated to a trust point, which means only one identity certificate is allowed from a CA.
- If multiple identity certificates (each from a distinct CA) have been obtained, the certificate that an application selects to use in a security protocol exchange with a peer is application specific.
- You do not need to designate one or more trust points for an application. Any application can use any certificate issued by any trust point as long as the certificate purpose satisfies the application requirements.
- You do not need more than one identity certificate from a trust point or more than one key-pair to be associated to a trust point. A CA certifies a given identity (name) only once and does not issue multiple certificates with the same subject name. If you need more than one identity certificate for a CA, then define another trust point for the same CA, associate another key-pair to it, and have it certified, provided CA allows multiple certificates with the same subject name.

Multiple Trusted CA Support

An MDS switch can be configured to trust multiple CAs by configuring multiple trust points and associating each with a distinct CA. With multiple trusted CAs, you do not have to enroll a switch with the specific CA that issued a certificate to a peer. Instead, you configure the switch with multiple trusted CAs that the peer trusts. A switch can then use a configured trusted CA to verify certificates offered by a peer that were not issued by the same CA defined in the identity of the switch.

Configuring multiple trusted CAs allows two or more switches enrolled under different domains (different CAs) to verify the identity of each other when using IKE to set up IPsec tunnels.

PKI Enrollment Support

Enrollment is the process of obtaining an identity certificate for the switch that is used for applications such as IPsec/IKE or SSH. It occurs between the switch requesting the certificate and the certificate authority.

The PKI enrollment process for a switch involves the following steps:

1. Generate an RSA private and public key-pair on the switch.
2. Generate a certificate request in standard format and forward it to the CA.
3. Manual intervention at the CA server by the CA administrator may be required to approve the enrollment request, when it is received by the CA.
4. Receive the issued certificate back from the CA, signed with the CA's private key.
5. Write the certificate into a nonvolatile storage area on the switch (bootflash).

Manual Enrollment Using Cut-and-Paste Method

Cisco MDS NX-OS supports certificate retrieval and enrollment using a manual cut-and-paste method. Cut-and-paste enrollment literally means you must cut and paste the certificate requests and resulting certificates between the switch and the CA, as follows:

1. Create an enrollment certificate request, which is displayed in base64-encoded text form.
2. Cut and paste the encoded certificate request text in an e-mail message or in a web form and send it to the CA.
3. Receive the issued certificate (in base64-encoded text form) from the CA in an e-mail message or in a web browser download.

4. Cut and paste the issued certificate to the switch using the certificate import facility.

**Note**

DCNM-SAN does not support cut and paste. Instead, it allows the enrollment request (certificate signing request) to be saved in a file to be sent manually to the CA.

Multiple RSA Key-Pair and Identity CA Support

Multiple identity CA support enables the switch to enroll with more than one trust point. This results in multiple identity certificates; each from a distinct CA. This allows the switch to participate in IPsec and other applications with many peers using certificates issued by appropriate CAs that are acceptable to those peers.

The multiple RSA key-pair support feature allows the switch to maintain a distinct key pair for each CA with which it is enrolled. Thus, it can match policy requirements for each CA without conflicting with the requirements specified by the other CAs, such as key length. The switch can generate multiple RSA key-pairs and associate each key-pair with a distinct trust point. Thereafter, when enrolling with a trust point, the associated key-pair is used to construct the certificate request.

Peer Certificate Verification

The PKI support on an MDS switch provides the means to verify peer certificates. The switch verifies certificates presented by peers during security exchanges pertaining to applications, such as IPsec/IKE and SSH. The applications verify the validity of the peer certificates presented to them. The peer certificate verification process involves the following steps:

- Verifies that the peer certificate is issued by one of the locally trusted CAs.
- Verifies that the peer certificate is valid (not expired) with respect to current time.
- Verifies that the peer certificate is not yet revoked by the issuing CA.

For revocation checking, two methods are supported: certificate revocation list (CRL) and Online Certificate Status Protocol (OCSP). A trust point uses one or both of these methods to verify that the peer certificate has not been revoked.

CRL Downloading, Caching, and Checking Support

Certificate revocation lists (CRLs) are maintained by CAs to give information of prematurely revoked certificates, and the CRLs are published in a repository. The download URL is made public and also specified in all issued certificates. A client verifying a peer's certificate should obtain the latest CRL from the issuing CA and use it to determine if the certificate has been revoked. A client can cache the CRLs of some or all of its trusted CAs locally and use them later if necessary until the CRLs expire.

Cisco MDS NX-OS allows the manual configuration of pre-downloaded CRLs for the trust points, and then caches them in the switch bootflash (cert-store). During the verification of a peer certificate by IPsec or SSH, the issuing CA's CRL is consulted only if the CRL has already been cached locally and the revocation checking is configured to use CRL. Otherwise, CRL checking is not performed and the certificate is considered to be not revoked if no other revocation checking methods are configured. This mode of CRL checking is called CRL optional.

OCSP Support

Online Certificate Status Protocol (OCSP) facilitates online certificate revocation checking. You can specify an OCSP URL for each trust point. Applications choose the revocation checking mechanisms in a specified order. The choices are CRL, OCSP, none, or a combination of these methods.

Import and Export Support for Certificates and Associated Key-Pairs

As part of the CA authentication and enrollment process, the subordinate CA certificate (or certificate chain) and identity certificates can be imported in standard PEM (base64) format.

The complete identity information in a trust point can be exported to a file in the password-protected PKCS#12 standard format. It can be later imported to the same switch (for example, after a system crash) or to a replacement switch. The information in a PKCS#12 file consists of the RSA key-pair, the identity certificate, and the CA certificate (or chain).

Maximum Limits

[Table 80: Maximum Limits for CA and Digital Certificate](#), on page 659 lists the maximum limits for CAs and digital certificate parameters.

Table 80: Maximum Limits for CA and Digital Certificate

Feature	Maximum Limit
Trust points declared on a switch	16
RSA key-pairs generated on a switch	16
Identity certificates configured on a switch	16
Certificates in a CA certificate chain	10
Trust points authenticated to a specific CA	10

Default Settings

[Table 81: Default CA and Digital Certificate Parameters](#), on page 659 lists the default settings for CAs and digital certificate parameters.

Table 81: Default CA and Digital Certificate Parameters

Parameters	Default
Trust point	None
RSA key-pair	None
RSA key-pair label	Switch FQDN
RSA key-pair modulus	512

Parameters	Default
RSA key-pair exportable	Yes
Revocation check method of trust point	CRL

Configuring CAs and Digital Certificates

This section describes the tasks you must perform to allow CAs and digital certificates your Cisco MDS switch device to interoperate. This section includes the following sections:

Generating an RSA Key Pair

RSA key-pairs are used to sign and/or encrypt and decrypt the security payload during security protocol exchanges for applications such as IKE/IPsec and SSH, and they are required before you can obtain a certificate for your switch.

To generate an RSA key-pair, follow these steps:

Procedure

-
- Step 1** Expand **Switches > Security** and then select **PKI** in the Information pane.
- Step 2** Click the **RSA Key-Pair** tab.
- Step 3** Click **Create Row**.
- Step 4** Select the switches for which you want to create the RSA key-pair.
- Step 5** Assign a name to the RSA key-pair.
- Step 6** Select the Size or modulus values. Valid modulus values are 512, 768, 1024, 1536, and 2048.
- Note** The security policy (or requirement) at the local site (MDS switch) and at the CA (where enrollment is planned) are considered in deciding the appropriate key modulus.
- Note** The maximum number of key-pairs you can configure on a switch is 16.
- Step 7** Check the **Exportable** check box if you want the key to be exportable.
- Caution** The exportability of a key-pair cannot be changed after key-pair generation.
- Note** Only exportable key-pairs can be exported in PKCS#12 format.
- Step 8** Click **Create** to create the RSA key pair.
-

Creating a Trust Point CA Association

To create a trust point CA association, follow these steps:

Procedure

-
- Step 1** Expand **Switches > Security**, and then select **PKI** in the Physical Attributes pane.

- Step 2** Click the **Trust Point** tab in the Information Pane.
- Step 3** Click **Create Row**.
- Step 4** Select the switch for which you are creating the trust point CA from the **Switch** drop-down menu.
- Step 5** Assign a name to the trust point CA.
- Step 6** Select a key-pair name to be associated with this trust point for enrollment. It was generated earlier in the [Generating an RSA Key Pair, on page 660](#). Only one RSA key-pair can be specified per CA.
- Step 7** From the RevokeCheckMethod drop-down menu, select the certificate revocation method that you would like to use. You can use CRL, OCSP, CRL OCSP, or OCSP CRL to check for certificate revocation.
- The CRL OCSP option checks for revoked certificates first in the locally stored CRL. If not found, the switch uses OCSP to check the revoked certificates on the URL specified in Step 7.
- Step 8** Enter the OCSP URL if you selected an OCSP certificate revocation method.
- Note** The OCSP URL must be configured before configuring the revocation checking method.
- Step 9** Click **Create** to successfully create the trust point CA.

Copying Files to Bootflash

To copy files to bootflash using Device Manager, follow these steps:

Procedure

- Step 1** Choose **Admin > Flash Files**.
- Step 2** Select bootflash in the Device field.
- Step 3** Click **Copy**.
- Step 4** Select **tftp** as the Protocol field.
- Step 5** Click the Browse button to locate the appropriate file to copy to bootflash.
- Step 6** Click **Apply** to apply these changes.

Authenticating the CA

The configuration process of trusting a CA is complete only when the CA is authenticated to the MDS switch. The switch must authenticate the CA. It does this by obtaining the self-signed certificate of the CA in PEM format, which contains the public key of the CA. Because the certificate of the CA is self-signed (the CA signs its own certificate) the public key of the CA should be manually authenticated by contacting the CA administrator to compare the fingerprint of the CA certificate.



Note

If the CA being authenticated is not a self-signed CA (that is, it is a subordinate CA to another CA, which itself may be a subordinate to yet another CA, and so on, finally ending in a self-signed CA), then the full list of the CA certificates of all the CAs in the certification chain needs to be input during the CA authentication step. This is called the *CA certificate chain* of the CA being authenticated. The maximum number of certificates in a CA certificate chain is 10.

To authenticate a CA, follow these steps:

Procedure

-
- Step 1** Expand **Switches > Security**, and then select **PKI** in the Physical Attributes pane.
- Step 2** Click the **Trust Point Actions** tab in the Information pane.
- Step 3** From the Command field drop-down menu, select the appropriate option.
- Available options are **caauth**, **cadelete**, **certreq**, **certimport**, **certdelete**, **pkcs12import**, and **pkcs12export**. The **caauth** option is provided to authenticate a CA and install its CA certificate or certificate chain in a trust point.
- Step 4** Click the Browse button in the URL field and select the appropriate import certificate file from the **Bootflash Files** dialog box. It is the file name containing the CA certificate or chain in the bootflash:filename format.
- Note** You can authenticate a maximum of 10 trust points to a specific CA.
- Note** If you do not see the required file in the Import Certificate dialog box, make sure that you copy the file to bootflash. See [Copying Files to Bootflash, on page 661](#).
- Step 5** Click **Apply Changes to save the changes**.
- Authentication is then confirmed or not confirmed depending on whether or not the certificate can be accepted after manual verification of its fingerprint.
-

Confirming CA Authentication

As mentioned in step 5 of [Authenticating the CA, on page 661](#), CA authentication is required to be followed by CA confirmation in order to accept the CA certificate based on its fingerprint verification.

To confirm CA authentication, follow these steps:

Procedure

-
- Step 1** Expand **Switches > Security**, and then select **PKI** in the Physical Attributes pane.
- Step 2** Click the **Trust Point Actions** tab in the Information Pane.
- Step 3** Make a note of the CA certificate fingerprint displayed in the IssuerCert FingerPrint column for the trust point row in question. Compare the CA certificate fingerprint with the fingerprint already communicated by the CA (obtained from the CA web site).
- If the fingerprints match exactly, accept the CA with the **certconfirm** command in the Command drop-down menu. Otherwise, reject the CA with the **certnoconfirm** command.
- Step 4** If you selected **certconfirm** in step 3, click Command and select the **certconfirm** action from the drop-down menu. Click **Apply Changes**.
- If you selected **certnoconfirm** in step 3, click Command and select the **certnoconfirm** action drop-down menu. Click **Apply Changes**.
-

Generating Certificate Requests

You must generate a request to obtain identity certificates from the associated trust point CA for each of your switch's RSA key pairs. You must then cut and paste the displayed request into an e-mail message or in a website form for the CA.

To generate a request for signed certificates from the CA, follow these steps:

Procedure

-
- Step 1** Expand **Switches > Security**, and then select **PKI** in the Physical Attributes pane.
- Step 2** Click the **Trust Point Actions** tab in the Information pane.
- Step 3** Select the **certreq** option from the Command drop-down menu.
- This action generates a PKCS#10 certificate signing request (CSR) needed for an identity certificate from the CA corresponding to this trust point entry. This entry requires an associated key-pair. The CA certificate or certificate chain should already be configured through the **caauth** action. See [Authenticating the CA, on page 661](#).
- Step 4** Enter the output file name for storing the generated certificate request.
- It will be used to store the CSR generated in PEM format. Use the format `bootflash:filename`. This CSR should be submitted to the CA to get the identity certificate. Once the identity certificate is obtained, it should be installed in this trust point. See [Installing Identity Certificates, on page 663](#).
- Step 5** Enter the *challenge* password to be included in the CSR.
- Note** The challenge password is not saved with the configuration. This password is required in the event that your certificate needs to be revoked, so you must remember this password.
- Step 6** Click **Apply Changes** to save the changes.
-

Installing Identity Certificates

You receive the identity certificate from the CA by e-mail or through a web browser in base64 encoded text form. You must install the identity certificate from the CA by cutting and pasting the encoded text using the CLI import facility.

To install an identity certificate received from the CA by e-mail or through a web browser, follow these steps:

Procedure

-
- Step 1** Expand **Switches > Security**, and then select **PKI** in the Physical Attributes pane.
- Step 2** Click the **Trust Point Actions** tab, in the Information pane.
- Step 3** Select the **certimport** option from the Command drop-down menu to import an identity certificate in this trust point. The identity certificate is obtained from the corresponding CA for a CSR that was generated previously (see [Generating Certificate Requests, on page 663](#)).
- Note** The identity certificate should be available in PEM format in a file in bootflash.

- Step 4** Enter the name of the certificate file that should have been copied to bootflash in the URL field in the bootflash:filename format.
- Step 5** Click **Apply Changes to save your changes**.
- If successful, the values of the identity certificate and its related objects, like the certificate file name, are automatically updated with the appropriate values as per the corresponding attributes in the identity certificate.

Saving Your Configuration

Save your work when you make configuration changes or the information is lost when you exit.

To save your configuration, follow these steps:

Procedure

- Step 1** Expand **Switches**, and then select **Copy Configuration** in the Physical Attributes pane.
- Step 2** Select the switch configuration including the RSA key pairs and certificates.
- Step 3** Click **Apply Changes** to save the changes.

Ensuring Trust Point Configurations Persist Across Reboots

The trust point configuration is a normal Cisco NX-OS configuration that persists across system reboots only if you copy it explicitly to the startup configuration. The certificates, key pairs, and CRL associated with a trust point are automatically persistent if you have already copied the trust point configuration in the startup configuration. Conversely, if the trust point configuration is not copied to the startup configuration, the certificates, key pairs, and CRL associated with it are not persistent since they require the corresponding trust point configuration after a reboot. Always copy the running configuration to the startup configuration to ensure that the configured certificates, key pairs, and CRLs are persistent. Also, save the running configuration after deleting a certificate or key pair to ensure that the deletions are permanent.

The certificates and CRL associated with a trust point automatically become persistent when imported (that is, without an explicitly copying to the startup configuration) if the specific trust point is already saved in startup configuration.

We also recommend that you create a password-protected backup of the identity certificates and save it to an external server (see *Exporting and Importing Identity Information in PKCS12 Format*).



Note Copying the configuration to an external server does include the certificates and key pairs.

Monitoring and Maintaining CA and Certificates Configuration

The tasks in the section are optional. This section includes the following topics:

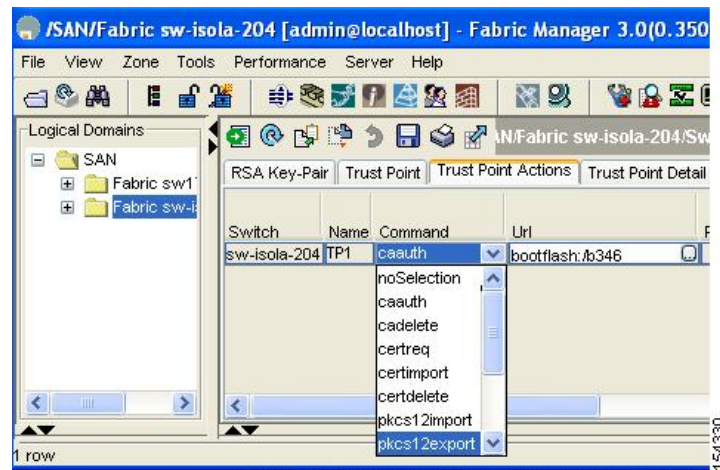
Exporting and Importing Identity Information in PKCS12 Format

To export a certificate and key pair to a PKCS#12-formatted file, follow these steps:

Procedure

- Step 1** Expand **Switches > Security**, and then select **PKI** in the Physical Attributes pane.
- Step 2** Click the **Trust Point Actions** tab in the Information Pane (see [Figure 92: Pkcs12export Option Exports a Key Pair](#), on page 665).
- Step 3** Select the **pkcs12export** option in the Command drop-down menu to export the key pair, identity certificate, and the CA certificate or certificate chain in PKCS#12 format from the selected trust point.

Figure 92: Pkcs12export Option Exports a Key Pair



- Step 4** Enter the output file name as bootflash:filename to store the exported PKCS#12 identity.
- Step 5** Enter the required password. The password is set for encoding the PKCS#12 data. On successful completion, the exported data is available in bootflash in the specified file.
- Step 6** Click **Apply Changes** to save the changes.

Importing Identity Information in PKCS12 Format

To import a certificate and key pair formatted as a PKCS#12 formatted file, follow these steps:

Procedure

- Step 1** Expand **Switches > Security**, and then select **PKI** in the Physical Attributes pane.
- Step 2** Click the **Trust Point Actions** tab in the Information pane (see [Figure 92: Pkcs12export Option Exports a Key Pair](#), on page 665).
- Step 3** Select the **pkcs12import** option from the Command drop-down menu to import the key-pair, identity certificate, and the CA certificate or certificate chain in the PKCS#12 format to the selected trust point.
- Step 4** Enter the input in the bootflash:filename format containing the PKCS#12 identity.

- Step 5** Enter the required password. The password is set for decoding the PKCS#12 data. On completion, the imported data is available in bootflash in the specified file.
- Step 6** Click **Apply Changes** to save the changes.
- On completion the trust point is created in the RSA key-pair table corresponding to the imported key pair. The certificate information is updated in the trust point.

Configuring a CRL

To import the CRL from a file to a trust point, follow these steps:

Procedure

- Step 1** Click **Switches > Security > PKI** in the Physical Attributes pane.
- Step 2** Click the **Trust Point Actions** tab in the Information pane.
- Step 3** Select the **crlimport** option from the Command drop-down menu to import the CRL to the selected trust point.
- Step 4** Enter the input file name with the CRL in the bootflash:filename format, in the URL field.
- Step 5** Click **Apply Changes** to save the changes.

Deleting Certificates from the CA Configuration

You can delete the identity certificates and CA certificates that are configured in a trust point. You must first delete the identity certificate, followed by the CA certificates. After deleting the identity certificate, you can disassociate the RSA key-pair from a trust point. The certificate deletion is necessary to remove expired or revoked certificates, certificates whose key-pairs are compromised (or suspected to be compromised) or CAs that are no longer trusted.

To delete the CA certificate (or the entire chain in the case of a subordinate CA) from a trust point, follow these steps:

Procedure

- Step 1** Click **Switches > Security > PKI** in the Physical Attributes pane.
- Step 2** Click the **Trust Point Actions** tab in the Information pane.
- Step 3** Select the **cadelete** option from the Command drop-down menu to delete the identity certificate from a trust point.
- Note** If the identity certificate being deleted is the last-most or only identity certificate in the device, you must use the **forcecertdelete** action to delete it. This ensures that the administrator does not mistakenly delete the last-most or only identity certificate and leave the applications (such as IKE and SSH) without a certificate to use.
- Step 4** Click **Apply Changes** to save the changes.

To delete the identity certificate, click the **Trust Point Actions** tab and select the **certdelete** or **forcecertdelete** in the Command drop-down menu.

Deleting RSA Key Pairs from Your Switch

Under certain circumstances you may want to delete your switch's RSA key pairs. For example, if you believe the RSA key pairs were compromised in some way and should no longer be used, you should delete the key pairs.

To delete RSA key-pairs from your switch, follow these steps:

Procedure

- Step 1** Expand **Switches > Security**, and then select **PKI** in the Physical Attributes pane.
- Step 2** Click the **RSA Key-Pair** tab in the Information pane.
- Step 3** Click **Delete Row**.
- Step 4** Click **Yes** or **No** in the Confirmation dialog box.

Note After you delete RSA key pairs from a switch, ask the CA administrator to revoke your switch's certificates at the CA. You must supply the challenge password you created when you originally requested the certificates. See [Generating Certificate Requests, on page 663](#).

Configuration Examples

To configure certificates on an MDS switch, follow these steps:

Procedure

- Step 1** Choose Switches and set the LogicalName field to configure the switch host name.
- Step 2** Choose Switches > Interfaces > Management > DNS and set the DefaultDomainName field to configure.
- Step 3** To create an RSA key-pair for the switch, follow these steps:
 - a) Choose Switches > Security > PKI and select the RSA Key-Pair tab.
 - b) Click Create Row and set the name and size field.
 - c) Check the Exportable check box and click Create.
- Step 4** To create a trust point and associate the RSA key-pairs with it, follow these steps:
 - a) Choose Switches > Security > PKI and select the Trustpoints tab.
 - b) Click Create Row and set the TrustPointName field.
 - c) Select the RSA key-pairs from the KeyPairName drop-down menu.
 - d) Select the certificates revocation method from the CARevoke drop-down menu.
 - e) Click Create.
- Step 5** Choose Switches > Copy Configuration and click Apply Changes to copy the running to startup configuration and save the trustpoint and key pair.

Step 6 Download the CA certificate from the CA that you want to add as the trustpoint CA.

Step 7 To authenticate the CA that you want to enroll to the trust point, follow these steps:

- a) Using Device Manager, choose Admin > Flash Files and select Copy and tftp copy the CA certificate to bootflash.
- b) Using DCNM-SAN, choose Switches > Security > PKI and select the TrustPoint Actions tab.
- c) Select cauth from the Command drop-down menu.
- d) Click ... in the URL field and select the CA certificate from bootflash.
- e) Click Apply Changes to authenticate the CA that you want to enroll to the trust point.
- f) Click the **Trust Point Actions** tab in the Information Pane.
- g) Make a note of the CA certificate fingerprint displayed in the IssuerCert FingerPrint column for the trust point row in question. Compare the CA certificate fingerprint with the fingerprint already communicated by the CA (obtained from the CA web site). If the fingerprints match exactly, accept the CA by performing the **certconfirm** trust point action. Otherwise, reject the CA by performing the **certnoconfirm** trust point action.
- h) If you select **certconfirm** in step g, click the Trust Point Actions tab, select **certconfirm** from the command drop-down menu, and then click **Apply Changes**.
- i) If you select **certnoconfirm** in step g, click the Trust Point Actions tab, select the **certnoconfirm** from the command drop-down menu and then click **Apply Changes**.

Step 8 To generate a certificate request for enrolling with that trust point, follow these steps:

- a) Click the **Trust Point Actions** tab in the Information pane.
- b) Select **certreq** from the Command drop-down menu. This generates a PKCS#10 certificate signing request (CSR) needed for an identity certificate from the CA corresponding to this trust point entry.
- c) Enter the output file name for storing the generated certificate request. It should be specified in the bootflash:filename format and will be used to store the CSR generated in PEM format.
- d) Enter the *challenge* password to be included in the CSR. The challenge password is not saved with the configuration. This password is required in the event that your certificate needs to be revoked, so you must remember this password.
- e) Click **Apply Changes** to save the changes.

Step 9 Request an identity certificate from the CA.

Note The CA may require manual verification before issuing the identity certificate.

Step 10 To import the identity certificate, follow these steps:

- a) Using Device Manager, choose Admin > Flash Files and select Copy and use TFTP to copy the CA certificate to bootflash.
- b) Using DCNM-SAN, choose Switches > Security > PKI and click the TrustPoint Actions tab.
- c) Select the **certimport** option from the Command drop-down menu to import an identity certificate in this trust point.

Note The identity certificate should be available in PEM format in a file in bootflash.

- d) Enter the name of the certificate file which was copied to bootflash, in the URL field in the bootflash:filename format.
- e) Click **Apply Changes to save your changes**.

If successful, the values of the identity certificate and its related objects, like the certificate file name, are automatically updated with the appropriate values as per the corresponding attributes in the identity certificate.

Downloading a CA Certificate

To download a CA certificate from the Microsoft Certificate Services web interface, follow these steps:

Procedure

- Step 1** Click the **Retrieve the CA certificate or certificate revocation task** radio button in the Microsoft Certificate Services web interface and click the **Next** button.
 - Step 2** Select the CA certificate file to download from the displayed list. Click the **Base 64 encoded** radio button, and choose the **Download CA certificate** link.
 - Step 3** Click the **Open** button in the File Download dialog box.
 - Step 4** Click the **Copy to File** button in the Certificate dialog box and click **OK**.
 - Step 5** Select the **Base-64 encoded X.509 (CER)** on the Certificate Export Wizard dialog box and click **Next**.
 - Step 6** Enter the destination file name in the File name: text box on the Certificate Export Wizard dialog box and click **Next**.
 - Step 7** Click the **Finish** button on the Certificate Export Wizard dialog box.
 - Step 8** Display the CA certificate stored in Base-64 (PEM) format using the Microsoft Windows **type** command.
-

Requesting an Identity Certificate

To request an identity certificate from a Microsoft Certificate server using a PKCS#10 certificate signing request (CRS), follow these steps:

Procedure

- Step 1** Click the Request an identity certificate radio button on the Microsoft Certificate Services web interface and click **Next**.
- Step 2** Click the **Advanced Request** radio button and click **Next**.
- Step 3** Click the **Submit a certificate request using a base64 encoded PKCS#10 file or a renewal request using a base64 encoded PKCS#7 file** radio button and click **Next**.
- Step 4** Paste the base64 PKCS#10 certificate request in the Saved Request text box and click **Next**.

The certificate request is copied from the MDS switch console (see the [Generating Certificate Requests, on page 663](#)).
- Step 5** Wait one or two days until the certificate is issued by the CA administrator.
- Step 6** The CA administrator approves the certificate request.
- Step 7** Click the **Check on a pending certificate** radio button on the Microsoft Certificate Services web interface and click **Next**.

- Step 8** Select the certificate request you want to check and click **Next**.
 - Step 9** Select **Base 64 encoded** and click the **Download CA certificate** link.
 - Step 10** Click **Open** on the File Download dialog box.
 - Step 11** Click the **Details** tab on the Certificate dialog and click the **Copy to File** button. Click the **Base-64 encoded X.509 (.CER)** radio button on the Certificate Export Wizard dialog box and click **Next**.
 - Step 12** Enter the destination file name in the File name: text box on the Certificate Export Wizard dialog box, then click **Next**.
 - Step 13** Click **Finish**.
 - Step 14** Display the identity certificate in base64-encoded format using the Microsoft Windows **type** command.
-

Revoking a Certificate

To revoke a certificate using the Microsoft CA administrator program, follow these steps:

Procedure

- Step 1** Click the **Issued Certificates** folder on the Certification Authority tree. From the list, right-click the certificate you want to revoke.
 - Step 2** Select **All Tasks > Revoke Certificate**.
 - Step 3** Select a reason for the revocation from the Reason code drop-down list, and click **Yes**.
 - Step 4** Click the **Revoked Certificates** folder to list and verify the certificate revocation.
-

Generating and Publishing the CRL

To generate and publish the CRL using the Microsoft CA administrator program, follow these steps:

Procedure

- Step 1** Select **Action > All Tasks > Publish** on the Certification Authority screen.
 - Step 2** Click **Yes** on the Certificate Revocation List dialog box to publish the latest CRL.
-

Downloading the CRL

To download the CRL from the Microsoft CA website, follow these steps:

Procedure

- Step 1** Click **Request the CA certificate or certificate revocation list** radio button on the Microsoft Certificate Services web interface and click **Next**.
- Step 2** Click the **Download latest certificate revocation list** link.

- Step 3** Click **Save** in the File Download dialog box.
 - Step 4** Enter the destination file name in the Save As dialog box and click **Save**.
 - Step 5** Display the CRL using the Microsoft Windows **type** command.
-

Importing the CRL

To import the CRL to the trust point corresponding to the CA, follow these steps:

Procedure

- Step 1** Click **Switches > Security > PKI** in the Physical Attributes pane.
 - Step 2** Click the **Trust Point Actions** tab in the Information pane.
 - Step 3** Select the **crlimport** option from the Command drop-down menu to import the CRL to the selected trust point.
 - Step 4** Enter the input file name with the CRL in the bootflash:filename format, in the URL field.
 - Step 5** Click **Apply Changes** to save the changes.
- The identity certificate for the switch that was revoked (serial number 0A338EA1000000000074) is listed at the end.
-



CHAPTER 31

Configuring FC-SP and DHCHAP

- [Configuring FC-SP and DHCHAP, on page 673](#)

Configuring FC-SP and DHCHAP

This chapter includes the following topics:

Information About Fabric Authentication

Fibre Channel Security Protocol (FC-SP) capabilities provide switch-switch and host-switch authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) is an FC-SP protocol that provides authentication between Cisco MDS 9000 Family switches and other devices. DHCHAP consists of the CHAP protocol combined with the Diffie-Hellman exchange.

To authenticate through VFC ports, FC-SP peers use the port VSAN for communication. Hence, the port VSAN needs to be the same and active on both the peers to send and receive authentication messages.

All switches in the Cisco MDS 9000 Family enable fabric-wide authentication from one switch to another switch, or from a switch to a host. These switch and host authentications are performed locally or remotely in each fabric. As storage islands are consolidated and migrated to enterprise-wide fabrics new security challenges arise. The approach of securing storage islands cannot always be guaranteed in enterprise-wide fabrics.

For example, in a campus environment with geographically distributed switches someone could maliciously interconnect incompatible switches or you could accidentally do so, resulting in Inter-Switch Link (ISL) isolation and link disruption. This need for physical security is addressed by switches in the Cisco MDS 9000 Family.



Note

Fibre Channel (FC) host bus adapters (HBAs) with appropriate firmware and drivers are required for host-switch authentication.

DHCHAP

DHCHAP is an authentication protocol that authenticates the devices connecting to a switch. Fibre Channel authentication allows only trusted devices to be added to a fabric, 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.

Configuring the DHCHAP feature requires the ENTERPRISE_PKG license (see the Cisco MDS 9000 Family NX-OS Licensing Guide).

DHCHAP Compatibility with Existing Cisco MDS Features

This section identifies the impact of configuring the DHCHAP feature along with existing Cisco MDS features:

- PortChannel interfaces—If DHCHAP is enabled for ports belonging to a PortChannel, DHCHAP authentication is performed at the physical interface level, not at the PortChannel level.
- FCIP interfaces—The DHCHAP protocol works with the FCIP interface just as it would with a physical interface.
- Port security or fabric binding—Fabric binding policies are enforced based on identities authenticated by DHCHAP.
- VSANs—DHCHAP authentication is not done on a per-VSAN basis.
- High availability—DHCHAP authentication works transparently with existing HA features.

About Enabling DHCHAP

By default, the DHCHAP feature is disabled in all switches in the Cisco MDS 9000 Family.

You must explicitly enable the DHCHAP feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.

About DHCHAP Authentication Modes

The DHCHAP authentication status for each interface depends on the configured DHCHAP port mode.

When the DHCHAP feature is enabled in a switch, each Fibre Channel interface or FCIP interface may be configured to be in one of four DHCHAP port modes:

- On—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software moves the link to an isolated state.
- Auto-Active—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software continues with the rest of the initialization sequence.
- Auto-Passive (default)—The switch does not initiate DHCHAP authentication, but participates in DHCHAP authentication if the connecting device initiates DHCHAP authentication.
- Off—The switch does not support DHCHAP authentication. Authentication messages sent to such ports return error messages to the initiating switch.



Note Whenever DHCHAP port mode is changed to a mode other than the Off mode, reauthentication is performed.

[#unique_1497 unique_1497_Connect_42_tab_1000587](#) identifies the switch-to-switch authentication behavior between two Cisco MDS switches in various modes.

Table 82: DHCHAP Authentication Status Between Two MDS Switches

Switch NDHCHAP 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.	
auto-Passive				
off	Link is brought down.	FC-SP authentication is <i>not</i> performed.		

About the DHCHAP Hash Algorithm

Cisco MDS switches support a default hash algorithm priority list of MD5 followed by SHA-1 for DHCHAP authentication.



Tip If you change the hash algorithm configuration, then change it globally for all switches in the fabric.



Caution RADIUS and TACACS+ protocols always use MD5 for CHAP authentication. Using SHA-1 as the hash algorithm may prevent RADIUS and TACACS+ usage—even if these AAA protocols are enabled for DHCHAP authentication.

About the DHCHAP Group Settings

All switches in the Cisco MDS Family support all DHCHAP groups specified in the standard: 0 (null DH group, which does not perform the Diffie-Hellman exchange), 1, 2, 3, or 4.



Tip If you change the DH group configuration, change it globally for all switches in the fabric.

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 approaches to manage passwords for all switches in the fabric that participate in DHCHAP.

- Approach 1—Use the same password for all switches in the fabric. This is the simplest approach. When you add a new switch, you use the same password to authenticate that switch in this fabric. It is also the most vulnerable approach if someone from the outside maliciously attempts to access any one switch in the fabric.
- Approach 2—Use a different password for each switch and maintain that password list in each switch in the fabric. When you add a new switch, you create a new password list and update all switches with the new list. Accessing one switch yields the password list for all switches in that fabric.
- Approach 3—Use different passwords for different switches in the fabric. When you add a new switch, multiple new passwords corresponding to each switch in the fabric must be generated and configured in each switch. Even if one switch is compromised, the password of other switches are still protected. This approach requires considerable password maintenance by the user.



Note All passwords are restricted to 64 alphanumeric characters and can be changed, but not deleted.



Tip We recommend using RADIUS or TACACS+ for fabrics with more than five switches. If you need to use a local password database, you can continue to do so using Approach 3 and using the Cisco MDS 9000 Family DCNM-SAN to manage the password database.

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.

About the DHCHAP Timeout Value

During the DHCHAP protocol exchange, if the MDS switch does not receive the expected DHCHAP message within a specified time interval, authentication failure is assumed. The time ranges from 20 (no authentication is performed) to 1000 seconds. The default is 30 seconds.

When changing the timeout value, consider the following factors:

- The existing RADIUS and TACACS+ timeout values.

The same value must also be configured on all switches in the fabric.

Enabling FC-SP on ISLs

There is an ISL pop-up menu in DCNM-SAN 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.



Note FC-SP DHCHAP mode requires port-flap on both end of the ISL.

Default Settings

[Table 83: Default Fabric Security Settings](#), on page 677 lists the default settings for all fabric security features in any switch.

Table 83: 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 DHCHAP

To configure DHCHAP authentication using the local password database, follow these steps:

Procedure

- 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.

Enabling DHCHAP

To enable DHCHAP for a Cisco MDS switch, follow these steps:

Procedure

- Step 1** Expand Switches, expand Security and then select FC-SP.

The Control tab is the default. You see the FC-SP enable state for all switches in the fabric.

- Step 2** Set the Command drop-down menu to enable for all switches that you want to enable FC-SP on.
 - Step 3** Click the Apply Changes icon to enable FC-SP and DHCHAP on the selected switches.
-

Configuring the DHCHAP Mode

To configure the DHCHAP mode for a particular interface, follow these steps:

Procedure

- Step 1** Expand Switches, expand Interfaces, and then select FC Physical.
You see the interface configuration in the Information pane.
 - Step 2** Click the FC-SP tab.
 - Step 3** Set the **Mode** drop-down menu to the DHCHAP authentication mode you want to configure for that interface.
 - Step 4** Click the Apply Changes icon to save these DHCHAP port mode settings.
-

Configuring the DHCHAP Hash Algorithm

To configure the hash algorithm, follow these steps:

Procedure

- Step 1** Choose Switches > Security, and then select FC-SP.
 - Step 2** Click the General/Password tab.
You see the DHCHAP general settings mode for each switch.
 - 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.
-

Configuring the DHCHAP Group Settings

To change the DH group settings, follow these steps:

Procedure

- Step 1** Expand Switches > Security, and then select 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.
-

Configuring DHCHAP Passwords for the Local Switch

To configure the DHCHAP password for the local switch, follow these steps:

Procedure

- Step 1** Expand Switches > Security, and then select FC-SP.
You see the FC-SP configuration in the Information pane.
- Step 2** Click the Local Passwords tab.
- 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.
- Step 6** Click Create to save the updated password.
-

Configuring DHCHAP Passwords for Remote Devices

To locally configure the remote DHCHAP password for another switch in the fabric, follow these steps:

Procedure

- Step 1** Right-click an ISL and select Enable FC-SP from the drop-down list (see [#unique_1511 unique_1511_Connect_42_fig_1001034](#)).
You see the Enable FC-SP dialog box.
- Step 2** Click Apply to save the updated password.
-

Configuring the DHCHAP Timeout Value

To configure the DHCHAP timeout value, follow these steps:

Procedure

- Step 1** Expand Switches > Security, and then select 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.

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 individually set authentication options. If authentication is not configured, local authentication is used by default.



CHAPTER 32

Configuring Cisco TrustSec Fibre Channel Link Encryption

- [Configuring Cisco TrustSec Fibre Channel Link Encryption, on page 681](#)

Configuring Cisco TrustSec Fibre Channel Link Encryption

This chapter provides an overview of the Cisco TrustSec Fibre Channel (FC) Link Encryption feature and describes how to configure and set up link-level encryption between switches.

This chapter includes the following topics:

Information About Cisco TrustSec FC Link Encryption

Cisco TrustSec FC Link Encryption is an extension of the Fibre Channel-Security Protocol (FC-SP) feature and uses the existing FC-SP architecture to provide integrity and confidentiality of transactions. Encryption is now added to the peer authentication capability to provide security and prevent unwanted traffic interception. Peer authentication is implemented according to the FC-SP standard using the Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) protocol.



Note

Cisco TrustSec FC Link Encryption is currently only supported between Cisco MDS switches. This feature is not supported when you downgrade to software versions which do not have the Encapsulating Security Protocol (ESP) support.

This section includes the following topics:

Supported Modules

The following modules are supported for the Cisco TrustSec FC Link Encryption feature:

- 1/2/4/8 Gbps 24-Port Fibre Channel switching module (DS-X9224-96K9)
- 1/2/4/8 Gbps 48-Port Fibre Channel switching module (DS-X9248-96K9)
- 1/2/4/8 Gbps 4/44-Port Fibre Channel switching module (DS-X9248-48K9)

Cisco TrustSec FC Link Encryption Terminology

The following Cisco TrustSec FC Link Encryption-related terms are used in this chapter:

- **Galois Counter Mode (GCM)**—A block cipher mode of operation providing confidentiality and data-origin authentication.
- **Galois Message Authentication Code (GMAC)**—A block cipher mode of operation providing only data-origin authentication. It is the authentication-only variant of GCM.
- **Security Association (SA)**—A connection that handles the security credentials and controls how they propagate between switches. The SA includes parameters such as salt and keys.
- **Key**—A 128-bit hexadecimal string that is used for frame encryption and decryption. The default value is zero.
- **Salt**—A 32-bit hexadecimal number that is used during encryption and decryption. The same salt must be configured on both sides of the connection to ensure proper communication. The default value is zero.
- **Security Parameters Index (SPI) number**—A 32-bit number that identifies the SA to be configured to the hardware. The range is from 256 to 4,294,967,295.

Support for AES Encryption

The Advanced Encryption Standard (AES) is the symmetric cipher algorithm that provides a high-level of security, and can accept different key sizes.

The Cisco TrustSec FC Link Encryption feature supports the 128-bit AES for security encryption and enables either AES-GCM or AES-GMAC for an interface. The AES-GCM mode provides encryption and authentication of the frames and AES-GMAC provides only the authentication of the frames that are being passed between the two peers.

Guidelines and Limitations

This section lists the guidelines for Cisco TrustSec FC Link Encryption:

- Ensure that Cisco TrustSec FC Link Encryption is enabled only between MDS switches. This feature is supported only on E-ports or the ISLs, and errors will result if non-MDS switches are used.
- Ensure that the peers in the connection have the same configurations. If there are differences in the configurations, a “port re-init limit exceeded” error message is displayed.
- Before applying the SA to the ingress and egress hardware of a switch interface, ensure that the interface is in the admin shut mode.

Configuring Cisco TrustSec Fibre Channel Link Encryption

By default, the FC-SP feature and the Cisco TrustSec FC Link Encryption feature are disabled in all switches in the Cisco MDS 9000 Family.

You must explicitly enable the FC-SP feature to access the configuration and verification commands for fabric authentication and encryption. When you disable this feature, all related configurations are automatically discarded.

Setting Up Security Association Parameters using DCNM-SAN

To perform encryption between the switches, a security association (SA) needs to be set up. An administrator manually configures the SA before the encryption can take place. The SA includes parameters such as keys and salt, that are required for encryption. You can set up to 2000 SAs in a switch. The **no fcsp esp sa**

spl_number command returns an error saying that the SA is in use if the specified SA is currently programmed to the ports.

To determine which ports are using the SA, use the **show running-config fcsp** command.



Note Cisco TrustSec FC Link Encryption is currently supported only on DHCHAP on and off modes.

To set up the SA parameters, such as keys and salt, using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand Switches > Security, and then select FC-SP (DHCHAP).
You see the FC-SP configuration in the Information pane.
- Step 2** Click the SA tab.
You see the SA parameters for each switch.
- Step 3** Click the Create Row icon.
You see the Create SA Parameters dialog box.
- Step 4** Select the switches on which you want to perform an encryption.
- Step 5** Select a value for the SP. The range is from 256 to 65536.
- Step 6** Enter a value for the salt. Alternatively, click Salt Generator to select a value.
- Step 7** Enter a value for the key. Alternatively, click Key Generator to select a value.
- Step 8** Click Create to save the changes.

Setting Up Security Association Parameters using Device Manager

To set up the SA parameters, such as keys and salt, using Device Manager, follow these steps:

Procedure

- Step 1** Choose Switches > Security, and then select FC-SP.
You see the FC-SP configuration dialog box.
- Step 2** Click SA tab.
You see the SA parameters for each switch.
- Step 3** Click Create to create new parameters.
You see the Create FC-SP SA dialog box.
- Step 4** Select a value for the SP. The range is from 256 to 65536.
- Step 5** Enter a value for the salt. Alternatively, click Salt Generator to select a value.

- Step 6** Enter a value for the key. Alternatively, click Key Generator to select a value.
- Step 7** Click Create to save the changes.
-

Setting Up Security Association Parameters using Device Manager

To set up the SA parameters, such as keys and salt, using Device Manager, follow these steps:

Procedure

- Step 1** Choose Switches > Security, and then select FC-SP.
You see the FC-SP configuration dialog box.
- Step 2** Click SA tab.
You see the SA parameters for each switch.
- Step 3** Click Create to create new parameters.
You see the Create FC-SP SA dialog box.
- Step 4** Select a value for the SP. The range is from 256 to 65536.
- Step 5** Enter a value for the salt. Alternatively, click Salt Generator to select a value
- Step 6** Enter a value for the key. Alternatively, click Key Generator to select a value.
- Step 7** Click Create to save the changes.
-

Configuring ESP Settings

Once the SA is created, you need to configure Encapsulating Security Protocol (ESP) on the ports. You should specify the egress and ingress ports for the encryption and decryption of packets between the network peers. The egress SA specifies which keys or parameters are to be used for encrypting the packets that leave the switch. The ingress SA specifies which keys or parameters are to be used to decrypt the packets entering that particular port.

Configuring ESP Modes

Configure the ESP settings for the ports as GCM to enable message authentication and encryption or as GMAC to enable message authentication.

The default ESP mode is AES-GCM.

This section covers the following topics:

Configuring AES-GMAC

To configure ESP settings, follow these steps:

Procedure

- Step 1** Expand Switches > Security, and then select FC-SP (DHCHAP).
You see the FC-SP configuration in the Information pane.
- Step 2** Click the ESP Interfaces tab.
You see the Interface details for each switch.
- Step 3** Click the Create Row icon.
You see the Create ESP Interfaces dialog box.
- Step 4** Select the switches on which you want to perform an encryption.
- Step 5** Enter an interface for the selected switch.
- Step 6** Select the appropriate ESP mode for the encryption.
- Step 7** Enter the appropriate egress port for the encryption.
- Step 8** Enter the appropriate ingress port for the encryption.
- Step 9** Click Create to save the changes.
-

Configuring AES-GMAC using Device Manager

To configure ESP settings using Device Manager, follow these steps:

Procedure

- Step 1** Expand Switches > Security, and then select FC-SP.
You see the FC-SP configuration dialog box.
- Step 2** Click the ESP Interfaces tab.
You see the Interface details for each switch.
- Step 3** Click Create.
You see the Create FC-SP ESP Interfaces dialog box.
- Step 4** Enter an interface for any switch for encryption. Alternatively, you can select values from the available interfaces for the selected switch.
- Step 5** Select the appropriate ESP mode for the encryption.
- Step 6** Enter the appropriate egress port for the encryption.
- Step 7** Enter the appropriate ingress port for the encryption.
- Step 8** Click Create to save the changes.
-

Configuring ESP Using ESP Wizard

You can configure and set up link-level encryption between switches using ESP wizard. You can configure an existing Inter-Switch Link (ISL) as a secure ISL or edit an existing secure ingress SPI and egress SPI using this wizard.

To configure ESP using ESP wizard, follow these steps:

Procedure

-
- Step 1** Right-click Tools > Security> FC-SP ESP Link Security to launch the ESP wizard from DCNM-SAN.
- Step 2** Select the appropriate ISL to secure or edit security.
- Note** Only ISLs with FC-SP port mode turned on and available on ESP- capable switches or blades are displayed.
- Step 3** Create new Security Associations (SAs).
- You can create a new SA for each switch or use the existing SAs. You can click View Existing SA to view the existing SAs.
- Note** The existing list of SAs displays all existing SAs for a switch. The wizard runs only when a pair of switches have a common SA. The wizard checks for this requirement when you select Next and a warning message is displayed if a pair of switches do not have a common SA. You must create a common SA on the pair of the switches to run this wizard.
- Step 4** Specify the Egress port, Ingress port, and ESP mode for the selected ISL.
- The Egress and Ingress ports are auto populated with SPIs of the SAs common to a pair of switches in case of a secured ISL.
- In this scenario, the mode is disabled and you cannot edit the modes for a secured ISL.
- Note** You can modify an existing ESP configuration provided the selected ISLs are enabled.
- Step 5** Review your configuration.
- Step 6** Click Finish to start the configuration for the ESP setup. You can view the status of the configuration in the status column.
-

Verifying Cisco TrustSec Fibre Channel Link Encryption Configuration

You can view information about the Cisco TrustSec FC Link Encryption feature using the show commands in DCNM-SAN or Device Manager.

Command	Purpose
show fcsp interface fc7/41	Displays all FC-SP-related information for a specific interface.
show running-config fcsp	Displays all the run-time information relevant to FC-SP.
show fcsp interface fc3/31 statistics	Displays all statistics related to DHCHAP and ESP for an interface.

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference*.

This section has the following topics:

Displaying FC-SP Interface Statistics

You can view the statistics data that displays the Encapsulating Security Protocol-ESP Security Parameter (SPI) mismatches and Interface-Encapsulating Security Protocol authentication failures information using DCNM-SAN.

To view the ESP statistics for an interface, follow these steps:

Procedure

- Step 1** Expand Interfaces > FC Physical, and then select FC-SP.
You see the FC-SP configuration in the Information pane.
 - Step 2** Click the FC-SP tab.
You see view the FC-SP statistics data in the Information pane.
 - Step 3** Click Refresh to refresh the statistics data.
-

Displaying FC-SP Interface Statistics Using Device Manager

To view the ESP statistics for an interface using Device Manager, follow these steps:

Procedure

- Step 1** Choose Security > FC Physical, and then select FC-SP.
You see the FC-SP configuration in the Information pane.
 - Step 2** Click the Statistics tab.
You see the statistics in the Information pane.
 - Step 3** Click Refresh to refresh the statistics data.
-



CHAPTER 33

Configuring FIPS

- [Configuring FIPS, on page 689](#)

Configuring FIPS

The Federal Information Processing Standards (FIPS) Publication 140-2, Security Requirements for Cryptographic Modules, details the U.S. government requirements for cryptographic modules. FIPS 140-2 specifies that a cryptographic module should be a set of hardware, software, firmware, or some combination that implements cryptographic functions or processes, including cryptographic algorithms and, optionally, key generation, and is contained within a defined cryptographic boundary.

FIPS specifies certain crypto algorithms as secure, and it also identifies which algorithms should be used if a cryptographic module is to be called FIPS compliant.



Note

Cisco MDS SAN-OS Release 3.1(1) and NX-OS Release 4.1(1b) or later implements FIPS features and is currently in the certification process with the U.S. government, but it is not FIPS compliant at this time.

This chapter includes the following topics:

Information About FIPS Self-Tests

A cryptographic module must perform power-up self-tests and conditional self-tests to ensure that it is functional.



Note

FIPS power-up self-tests automatically run when FIPS mode is enabled by entering the `fips mode enable` command. A switch is in FIPS mode only after all self-tests are successfully completed. If any of the self-tests fail, then the switch is rebooted.

Power-up self-tests run immediately after FIPS mode is enabled. A cryptographic algorithm test using a known answer must be run for all cryptographic functions for each FIPS 140-2-approved cryptographic algorithm implemented on the Cisco MDS 9000 Family.

Using a known-answer test (KAT), a cryptographic algorithm is run on data for which the correct output is already known, and then the calculated output is compared to the previously generated output. If the calculated output does not equal the known answer, the known-answer test fails.

Conditional self-tests must be run when an applicable security function or operation is invoked. Unlike the power-up self-tests, conditional self-tests are executed each time their associated function is accessed.

Conditional self-tests include the following:

- Pair-wise consistency test—This test is run when a public-private key-pair is generated.
- Continuous random number generator test—This test is run when a random number is generated.

Both of these tests automatically run when a switch is in FIPS mode.

Guidelines and Limitations

Follow these guidelines before enabling FIPS mode:

- Make your passwords a minimum of eight characters in length.
- Disable Telnet. Users should log in using SSH only.
- Disable remote authentication through RADIUS/TACACS+. Only users local to the switch can be authenticated.
- Disable SNMP v1 and v2. Any existing user accounts on the switch that have been configured for SNMPv3 should be configured only with SHA for authentication and AES/3DES for privacy.
- Disable VRRP.
- Delete all IKE policies that either have MD5 for authentication or DES for encryption. Modify the policies so they use SHA for authentication and 3DES/AES for encryption.
- Delete all SSH Server RSA1 key-pairs.

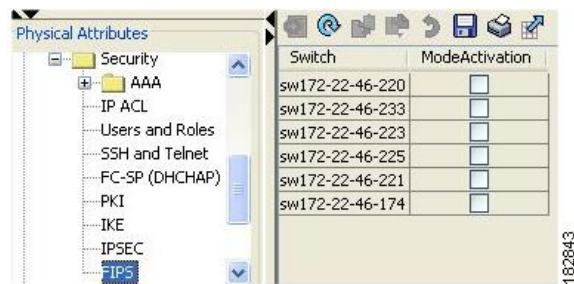
Enabling FIPS Mode using DCNM-SAN

To enable FIPS mode using DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Expand Switches from the Physical Attributes pane. Expand Security and then select FIPS. You see the FIPS activation details in the Information pane.

Figure 93: FIPS Activation in DCNM-SAN



- Step 2** Check the ModeActivation check box next to the switch for which you want to enable FIPS mode.
- Step 3** Click Apply Changes to commit and distribute these changes.
- Step 4** Click Undo Changes to discard any unsaved changes.

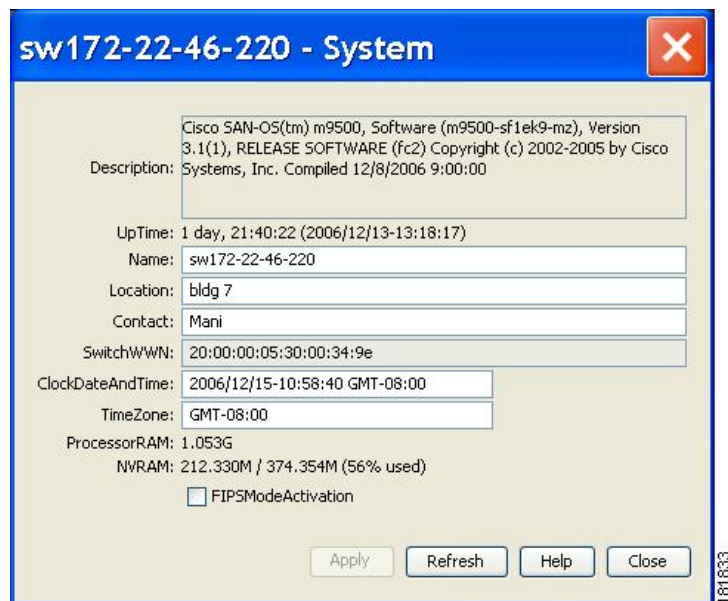
Enabling FIPS Mode using DCNM Manager

To enable FIPS mode using Device Manager, follow these steps:

Procedure

- Step 1** Choose Physical > System or right-click and select Configure. You see the System dialog box.

Figure 94: System Dialog Box



- Step 2** Check the FIPSMODEActivation check box to enable FIPS mode on the selected switch.

Step 3 Click Apply to save the changes.

Step 4 Click Close to close the dialog box.

Field Descriptions for FIPS

FIPS

Field	Description
ModeActivation	<p>To enable/disable FIPS mode on the device. FIPS 140-2 is a set of security requirements for cryptographic modules and it details the U.S. Government requirements for cryptographic modules. A module will comprise both hardware and software, eg a datacenter switching or routing module.</p> <p>The module is said to be in FIPS enabled mode when a request is recieved to enable the FIPS mode and a set of self-tests are successfully run in response to the request. If the self-tests fail, then an appropriate error is returned.</p>



CHAPTER 34

Configuring IPv4 and IPv6 Access Control Lists

- [Configuring IPv4 and IPv6 Access Control Lists, on page 693](#)

Configuring IPv4 and IPv6 Access Control Lists

Cisco MDS 9000 Family switches can route IP version 4 (IPv4) traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature routes traffic between VSANs. To do so, each VSAN must be in a different IPv4 subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMS):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding on the in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

Switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. VRRP is a restartable application that provides a redundant, alternate path to the gateway switch.

IPv4 Access Control Lists (IPv4-ACLs and IPv6-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IPv4-ACLs and IPv6-ACLs restrict IP-related traffic based on the configured IP filters. A filter contains the rules to match an IP packet, and if the packet matches, the rule also stipulates if the packet should be permitted or denied.

Each switch in the Cisco MDS 9000 Family can have a maximum total of 128 IPv4-ACLs or 128 IPv6-ACLs and each IPv4-ACL or IPv6-ACL can have a maximum of 256 filters.

This chapter includes the following topics:

Information About IPv4 and IPv6 Access Control Lists

Cisco MDS 9000 Family switches can route IP version 4 (IPv4) traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature routes traffic between VSANs. To do so, each VSAN must be in a different IPv4 subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMS):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.

- IP forwarding on the in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

IPv4 Access Control Lists (IPv4-ACLs and IPv6-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IPv4-ACLs and IPv6-ACLs restrict IP-related traffic based on the configured IP filters. A filter contains the rules to match an IP packet, and if the packet matches, the rule also stipulates if the packet should be permitted or denied.

Each switch in the Cisco MDS 9000 Family can have a maximum total of 128 IPv4-ACLs or 128 IPv6-ACLs and each IPv4-ACL or IPv6-ACL can have a maximum of 256 filters.

This section contains the following topics:

About Filter Contents

An IP filter contains rules for matching an IP packet based on the protocol, address, port, ICMP type, and type of service (ToS).

Protocol Information

The protocol information is required in each filter. It identifies the name or number of an IP protocol. You can specify the IP protocol in one of two ways:

- Specify an integer ranging from 0 to 255. This number represents the IP protocol.
- Specify the name of a protocol including, but not restricted to, Internet Protocol (IP), Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Control Message Protocol (ICMP).



Note

When configuring IPv4-ACLs or IPv6-ACLs on Gigabit Ethernet interfaces, only use the TCP or ICMP options.

Address Information

The address information is required in each filter. It identifies the following details:

- Source—The address of the network or host from which the packet is being sent.
- Source-wildcard—The wildcard bits applied to the source.
- Destination—The number of the network or host to which the packet is being sent.
- Destination-wildcard—The wildcard bits applied to the destination.

Specify the source and source-wildcard or the destination and destination-wildcard in one of two ways:

- Using the 32-bit quantity in four-part, dotted decimal format (10.1.1.2/0.0.0.0 is the same as host 10.1.1.2).
 - Each wildcard bit set to zero indicates that the corresponding bit position in the packet's IPv4 address must exactly match the bit value in the corresponding bit position in the source.
 - Each wildcard bit set to one indicates that both a zero bit and a one bit in the corresponding position of the packet's IPv4 or IPv6 address will be considered a match to this access list entry. Place ones in the bit positions you want to ignore. For example, 0.0.255.255 requires an exact match of only

the first 16 bits of the source. Wildcard bits set to one do not need to be contiguous in the source-wildcard. For example, a source-wildcard of 0.255.0.64 would be valid.

- Using the **any** option as an abbreviation for a source and source-wildcard or destination and destination-wildcard (0.0.0.0/255.255.255.255)

Port Information

The port information is optional. To compare the source and destination ports, use the **eq** (equal) option, the **gt** (greater than) option, the **lt** (less than) option, or the **range** (range of ports) option. You can specify the port information in one of two ways:

- Specify the number of the port. Port numbers range from 0 to 65535. [Table 84: TCP and UDP Port Numbers](#), on page 695 displays the port numbers recognized by the Cisco NX-OS software for associated TCP and UDP ports.
- Specify the name of a TCP or UDP port as follows:
 - TCP port names can only be used when filtering TCP.
 - UDP port names can only be used when filtering UDP.

Table 84: TCP and UDP Port Numbers

Protocol	Port	Number
UDP	dns	53
	tftp	69
	ntp	123
	radius accounting	1646 or 1813
	radius authentication	1645 or 1812
	snmp	161
	snmp-trap	162
	syslog	514
TCP ³²	ftp	20
ftp-data	21	
ssh	22	
telnet	23	
smtp	25	
tasacs-ds	65	
www	80	

Protocol	Port	Number
sftp	115	
http	143	
wbem-http	5988	
wbem-https	5989	

³² If the TCP connection is already established, use the established option to find matches. A match occurs if the TCP datagram has the ACK, FIN, PSH, RST, or URG control bit set.

ICMP Information

IP packets can be filtered based on the following optional ICMP conditions:

- icmp-type—The ICMP message type is a number from 0 to 255.
- icmp-code—The ICMP message code is a number from 0 to 255.

[Table 85: ICMP Type Value](#), on page 696 displays the value for each ICMP type.

Table 85: ICMP Type Value

ICMP Type ³³	Code
echo	8
echo-reply	0
destination unreachable	3
traceroute	30
time exceeded	11

³³ ICMP redirect packets are always rejected.

ToS Information

IP packets can be filtered based on the following optional ToS conditions:

- ToS level—The level is specified by a number from 0 to 15.
- ToS name—The name can be max-reliability, max-throughput, min-delay, min-monetary-cost, and normal.

Guidelines and Limitations

Follow these guidelines when configuring IPv4-ACLs or IPv6-ACLs in any switch or director in the Cisco MDS 9000 Family:

- You can apply IPv4-ACLs or IPv6-ACLs to VSAN interfaces, the management interface, Gigabit Ethernet interfaces on IPS modules and MPS-14/2 modules, and Ethernet PortChannel interfaces.

**Tip**

If IPv4-ACLs or IPv6-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group. See the Cisco MDS 9000 Family NX-OS IP Services Configuration Guide IP Services Configuration Guide, Cisco DCNM for SAN for guidelines on configuring IPv4-ACLs.

**Caution**

Do not apply IPv4-ACLs or IPv6-ACLs to only one member of a PortChannel group. Apply IPv4-ACLs or IPv6-ACLs to the entire channel group.

- Configure the order of conditions accurately. As the IPv4-ACL or the IPv6-ACL filters are sequentially applied to the IP flows, only the first match determines the action taken. Subsequent matches are not considered. Be sure to configure the most important condition first. If no conditions match, the software drops the packet.
- Configure explicit deny on the IP Storage Gigabit Ethernet ports to apply IP ACLs because implicit deny does not take effect on these ports.

Configuring IPv4-ACLs or IPv6-ACLs

This section contains the following topics:

Creating IPv4-ACLs or IPv6-ACLs

Traffic coming into the switch is compared to IPv4-ACL or IPv6-ACL filters based on the order that the filters occur in the switch. New filters are added to the end of the IPv4-ACL or the IPv6-ACL. The switch keeps looking until it has a match. If no matches are found when the switch reaches the end of the filter, the traffic is denied. For this reason, you should have the frequently hit filters at the top of the filter. There is an *implied deny* for traffic that is not permitted. A single-entry IPv4-ACL or IPv6-ACL with only one deny entry has the effect of denying all traffic.

To configure an IPv4-ACL or an IPv6-ACL, follow these steps:

1. Create an IPv4-ACL or an IPv6-ACL by specifying a filter name and one or more access condition(s). Filters require the source and destination address to match a condition. Use optional keywords to configure finer granularity.

**Note**

The filter entries are executed in sequential order. You can only add the entries to the end of the list. Take care to add the entries in the correct order.

2. Apply the access filter to specified interfaces.

To create an ordered list of IP filters in a named IPv4-ACL or IPv6-ACL profile using the IPv4-ACL Wizard, follow these steps:

Procedure

Step 1 Click the **IP ACL Wizard** icon from the DCNM-SAN toolbar.

You see the IP ACL Wizard.



Step 2 Enter a name for the IP-ACL.

If you are creating an IPv6-ACL, check the IPv6 check box.

Step 3 Click **Add** to add a new rule to this IP-ACL. You see a new rule in the table with default values.

Step 4 Modify the Source IP and Source Mask as necessary for your filter.

Note The IP-ACL Wizard only creates inbound IP filters.

Step 5 Choose the appropriate filter type from the Application drop-down list.

Step 6 Choose **permit** or **deny** from the Action drop-down list.

Step 7 Repeat Step 3 through Step 6 for additional IP filters.

Step 8 Click **Up** or **Down** to order the filters in this IP-ACL.

Note Order the IP filters carefully. Traffic is compared to the IP filters in order. The first match is applied and the rest are ignored.

Step 9 Click **Next**.

Note You see a list of switches that you can apply this IP-ACL.

Step 10 Uncheck any switches that you do not want to apply this IP-ACL.

Step 11 Select the **Interface** you want to apply this IP-ACL.

Step 12 Click **Finish** to create this IP-ACL and apply it to the selected switches.

Creating IPv4-ACLs or IPv6-ACLs

To add entries to an existing IPv4-ACL or an IPv6-ACL using Device Manager, follow these steps:

Procedure

Step 1 Choose **Security > IP ACL**.

Step 2 Click **Create** to create an IP-ACL profile.

You see the Create IP ACL Profiles dialog box. Enter an IP-ACL profile name.

Step 3 Click **Create** and then click **Close**.

This creates a new IP-ACL profile.

- Step 4** Click the IP-ACL you created and click **Rules**.
- After you create an IPv4-ACL or an IPv6-ACL, you can add subsequent IP filters at the end of the IPv4-ACL or the IPv6-ACL if you are using Device Manager. DCNM-SAN allows you to reorder existing rules for a profile. You cannot insert filters in the middle of an IPv4-ACL or an IPv6-ACL. Each configured entry is automatically added to the end of a IPv4-ACL or an IPv6-ACL.
- Step 5** Click **Create** to create an IP filter.
- Step 6** Choose either **permit** or **deny** for the Action and set the IP Number in the Protocol field. The drop-down menu provides common filtered protocols.
- Step 7** Set the source IP address you want this filter to match against and the wildcard mask, or check the **any** check box to match this filter against any IP address.
- This creates an IP filter that will check the source IP address of frames.
- Note** The wildcard mask denotes a subset of the IP address you want to match against. This allows a range of addresses to match against this filter.
- Step 8** Set the transport layer source port range if the protocol chosen is TCP or UDP.
- Step 9** Repeat Step 7 and Step 8 for the destination IP address and port range.
- This creates an IP filter that will check the destination IP address of frames.
- Step 10** Set the ToS, ICMPType, and ICMPCode fields as appropriate.
- Step 11** Check the **TCPEstablished** check box if you want to match TCP connections with ACK,FIN,PSH,RST,SYN or URG control bits set.
- Step 12** Check the **LogEnabled** check box if you want to log all frames that match this IP filter.
- Step 13** Click **Create** to create this IP filter and add it to your IP-ACL.

Deleting IP-ACLs

Before you begin

You must delete the association between the IP-ACL and interfaces before deleting the IP-ACL.

To delete an IP-ACL, follow these steps:

Procedure

- Step 1** Expand **Switches > Security**, and then select **IP ACL** from the Physical Attributes pane.
- You see the IP-ACL configuration in the Information pane.
- Step 2** Click the **Profiles** tab.
- You see a list of switches, ACLs, and profile names.
- Step 3** Select the row you want to delete. To delete multiple rows, hold down the Shift key while selecting rows.

Step 4 Click **Delete Row**. The IP-ACLs are deleted.

Reading the IP-ACL Log Dump

Use the LogEnabled check box option during IP filter creation to log information about packets that match this filter. The log output displays the ACL number, permit or deny status, and port information.

Use the **log-deny** option at the end of a filter condition to log information about packets that match dropped entries. The log output displays the ACL number, permit or deny status, and port information.



Note

To capture these messages in a logging destination, you must configure severity level 7 for the kernel and ipacl facilities and severity level 7 for the logging destination: logfile, monitor or console. For example:
 switch# **config t**switch(config)# **logging level kernel 7**switch(config)# **logging level ipacl 7**switch(config)# **logging logfile message 7**

For the input ACL, the log displays the raw MAC information. The keyword “MAC=” does not refer to showing an Ethernet MAC frame with MAC address information. It refers to the Layer 2 MAC-layer information dumped to the log. For the output ACL, the raw Layer 2 information is not logged.

The following example is an input ACL log dump:

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN=vsan1 OUT=
MAC=10:00:00:05:30:00:47:df:10:00:00:05:30:00:8a:1f:aa:aa:03:00:00:08:00:45:00:00:54:00:00:40:00:40:01:0e:86:0b:
0b:0b:0c:0b:0b:02:08:00:ff:9c:01:15:05:00:6f:09:17:3f:80:02:01:00:08:09:0a:0b:0c:0d:0e:0f:10:11:12:13:14:15:16:17:18:19:1a:1b:1c:1d:1e:1f:20:
21:22:23:24:25:26:27:28:29:2a:2b SRC=11.11.11.12 DST=11.11.11.2 LEN=84 TOS=0x00 PREC=0x00
TTL=64 ID=0 DF PROTO=ICMP TYPE=8 CODE=0 ID=277 SEQ=1280
```

The following example is an output ACL log dump:

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN= OUT=vsan1 SRC=11.11.11.2 DST=11.11.11.12 LEN=84 TOS=0x00 PREC=0x00 TTL=255
ID=38095 PROTO=ICMP TYPE=0 CODE=0 ID=277 SEQ=1280
```

Applying an IP-ACL to an Interface


You can define IP-ACLs without applying them. However, the IP-ACLs will have no effect until they are applied to an interface on the switch. You can apply IP-ACLs to VSAN interfaces, the management interface, Gigabit Ethernet interfaces on IPS modules and MPS-14/2 modules, and Ethernet PortChannel interfaces.



Tip

Apply the IP-ACL on the interface closest to the source of the traffic.

When you are trying to block traffic from source to destination, you can apply an inbound IPv4-ACL to M0 on Switch 1 instead of an outbound filter to M1 on Switch 3.



The **access-group** option controls access to an interface. Each interface can only be associated with one IP-ACL per direction. The ingress direction can have a different IP-ACL than the egress direction. The IP-ACL becomes active when applied to the interface.



Tip Create all conditions in an IP-ACL before applying it to the interface.



Caution If you apply an IP-ACL to an interface before creating it, all packets in that interface are dropped because the IP-ACL is empty.

The terms *in*, *out*, *source*, and *destination* are used as referenced by the switch:

- **In**—Traffic that arrives at the interface and goes through the switch; the source is where it transmitted from and the destination is where it is transmitted to (on the other side of the router).



Tip The IP-ACL applied to the interface for the ingress traffic affects both local and remote traffic.

- **Out**—Traffic that has already been through the switch and is leaving the interface; the source is where it transmitted from and the destination is where it is transmitted to.



Tip The IP-ACL applied to the interface for the egress traffic only affects local traffic.

Applying an IP-ACL to mgmt0

A system default ACL called mgmt0 exists on the mgmt0 interface. This ACL is not visible to the user, so mgmt0 is a reserved ACL name that cannot be used. The mgmt0 ACL blocks most ports and only allows access to required ports in compliance to accepted security policies.

To apply an IP-ACL to an interface, follow these steps:

Procedure

- Step 1** Expand **Switches > Security**, and then select **IP ACL** in the Physical Attributes pane.
You see the IP-ACL configuration in the Information pane.
- Step 2** Click the **Interfaces** tab.
You see a list of interfaces and associated IP-ACLs.
- Step 3** Click **Create Row**.

- Step 4** (Optional) Remove the switches you do not want to include in the IP-ACL by unchecking the check boxes next to the switch addresses.
- Set the **interface** you want associated with an IPv4-ACL or IPv6-ACL in the Interface field.
- Step 5** Choose a ProfileDirection (either **inbound** or **outbound**).
- Step 6** Enter the IP-ACL name in the Profile Name field.
- Note** This IP-ACL name must have already been created using the Create Profiles dialog box. If not, no filters will be enabled until you go to the Create Profiles dialog box and create the profile.
- Step 7** Click **Create** to associate the IP-ACL.
- You see the newly associated access list in the list of IP-ACLs.

Configuration Examples for IP-ACL

To define an IP-ACL that restricts management access using Device Manager, follow these steps:

Procedure

- Step 1** Choose **Security > IP ACL**.
- You see the IP-ACL dialog box.
- Step 2** Click **Create** to create an IP-ACL.
- You see the Create IP ACL Profiles dialog box.
- Step 3** Enter **RestrictMgmt** as the profile name and click **Create**.
- This creates an empty IP-ACL named RestrictMgmt.



- Step 4** Select **RestrictMgmt** and click **Rules**.
- You see an empty list of IP filters associated with this IP-ACL.
- Step 5** Click **Create** to create the first IP filter.
- You see the Create IP Filter dialog box.
- Step 6** Create an IP filter to allow management communications from a trusted subnet:

- a) Choose the **permit** Action and select **0 IP** from the Protocol drop-down menu.
- b) Set the source IP address to 10.67.16.0 and the wildcard mask to 0.0.0.255.

Note The wildcard mask denotes a subset of the IP address you want to match against. This allows a range of addresses to match against this filter.

- c) Check the **any** check box for the destination address.
- d) Click **Create** to create this IP filter and add it to the RestrictMgmt IP-ACL.

Step 7 Create an IP filter to allow ICMP ping commands:

- a) Choose the **permit** Action and select **1-ICMP** from the Protocol drop-down menu.
- b) Check the **any** check box for the source address.
- c) Check the **any** check box for the destination address.
- d) Select **8 echo** from the ICMPType drop-down menu.
- e) Click **Create** to create this IP filter and add it to the RestrictMgmt IP-ACL.

Repeat Step a through Step d to create an IP filter that blocks all other traffic.

Step 8 Create a final IP Filter to block all other traffic:

- a) Choose the **deny** Action and select **0 IP** from the Protocol drop-down menu.
- b) Check the **any** check box for the source address.
- c) Check the **any** check box for the destination address.
- d) Click **Create** to create this IP filter and add it to the RestrictMgmt IP-ACL.
- e) Click **Close** to close the Create IP Filter dialog box.

Repeat Step a through Step d to create an IP filter that blocks all other traffic.

Step 9 Apply the RestrictMgmt IP ACL to the mgmt0 interface:

- a) Click **Security**, select **IP ACL**, and then click the **Interfaces** tab in the IP ACL dialog box.
- b) Click **Create**.

You see the Create IP-ACL Interfaces dialog box.

- c) Select **mgmt0** from the Interfaces drop-down menu.
- d) Select the **inbound** Profile Director.
- e) Select **RestrictMgmt** from the ProfileName drop-down menu.
- f) Click **Create** to apply the RestrictMgmt IP-ACL to the mgmt0 interface.

Field Descriptions for IPv4 and IPv6 Access Control Lists

The following are the field descriptions for IPv4 and IPv6 access control lists:

IP ACL Profiles

Field	Description
Name	This is the unique IP protocol filter profile identifier.
Type	This object determines the usage type for this filter profile. This usage type cannot be changed after the profile has been created.

IP ACL Interfaces

Field	Description
ProfileName	This is the unique IP protocol filter profile identifier.

IP Filter Profiles

Field	Description
Action	If it is set to deny, all frames matching this filter will be discarded and scanning of the remainder of the filter list will be aborted. If it is set to permit, all frames matching this filter will be allowed for further bridging or routing processing.
Protocol	This filter protocol value matches the Internet Protocol Number in the frames. These IP numbers are defined in the Network Working Group Request for Comments (RFC) documents. Setting this to '-1' will make the filtering match any IP number.
Address	The source IP address to be matched for this filter. A value of 0 causes all source address to match.
Mask	This is the wildcard mask for the SrcAddress bits that must match. 0 bits in the mask indicate the corresponding bits in the SrcAddress must match in order for the matching to be successful, and 1 bits are don't care bits in the matching. A value of 0 causes only IP frames of source address the same as SrcAddress to match.
PortLow	If Protocol is UDP or TCP, this is the inclusive lower bound of the transport-layer source port range that is to be matched, otherwise it is ignored during matching. This value must be equal to or less than the value specified for this entry in SrcPortHigh.
PortHigh	If Protocol is UDP or TCP, this is the inclusive upper bound of the transport-layer source port range that is to be matched, otherwise it is ignored during matching. This value must be equal to or greater than the value specified for this entry in SrcPortLow. If this value is '0', the UDP or TCP port number is ignored during matching.
Address	The destination IP address to be matched for this filter. A value of 0 causes all source address to match.
Mask	This is the wildcard mask for the DestAddress bits that must match. 0 bits in the mask indicate the corresponding bits in the DestAddress must match in order for the matching to be successful, and 1 bits are don't care bits in the matching. A value of 0 causes only IP frames of source address the same as SrcAddress to match.
PortLow	If Protocol is UDP or TCP, this is the inclusive lower bound of the transport-layer destination port range that is to be matched, otherwise it is ignored during matching. This value must be equal to or less than the value specified for this entry in PortHigh.
PortHigh	If Protocol is UDP or TCP, this is the inclusive upper bound of the transport-layer destination port range that is to be matched, otherwise it is ignored during matching. This value must be equal to or greater than the value specified for this entry in DestPortLow. If this value is '0', the UDP or TCP port number is ignored during matching.

Field	Description
Precedence	<p>The IP traffic precedence parameters in each frame are used to guide the selection of the actual service parameters when transmitting a datagram through a particular network. Most network treats high precedence traffic as more important than other traffic. The IP Precedence value ranges from '0' to '7', with '7' the highest precedence and '0' the lowest precedence. The value '-1' means to match frames of any IP precedence. In other words, the IP precedence parameter will not be checked if this value is '-1'. The precedence level are:</p> <ul style="list-style-type: none"> • routine(0) - Routine traffic precedence • priority(1) - Priority traffic precedence • immediate(2) - Immediate traffic precedence • flash(3) - Flash traffic precedence • flashOverride(4) - Flash-override traffic precedence • critical(5) - Critical precedence • internet(6) - Internetwork control traffic precedence • network(7) - Network control traffic precedence.
TOS	The Type of Service (TOS) of the frame. The TOS values ranges from '0' to '15'. The value '-1' matches any TOS value.
ICMPType	This filter specifies the ICMP message type to be matched. Setting this value to '-1' will make the filtering match any ICMP message type.
ICMPCode	This filter specifies the ICMP message code to be matched. Setting this value to '-1' will make the filtering match any ICMP code.
TCPEstablished	This filter if true specifies that for TCP protocol, in an established connection, a match occurs if the TCP datagram has the ACK,FIN,PSH,RST,SYN or URG control bits set. If false, a match will occur for any TCP datagram.
LogEnabled	Specifies whether filtered frames will be logged by the filtering subsystem or not. If true, then all frames will be logged. If false, then no frame will be logged.



CHAPTER 35

Configuring IPsec Network Security

- [Configuring IPsec Network Security, on page 707](#)

Configuring IPsec Network Security

IP security (IPsec) protocol is a framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. It is developed by the Internet Engineering Task Force (IETF). IPsec provides security services at the IP layer, including protecting one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host. The overall IPsec implementation is the latest version of RFC 2401. Cisco NX-OS IPsec implements RFC 2402 through RFC 2410.

IPsec uses the Internet Key Exchange (IKE) protocol to handle protocol and algorithm negotiation and to generate the encryption and authentication keys used by IPsec. While IKE can be used with other protocols, its initial implementation is with the IPsec protocol. IKE provides authentication of the IPsec peers, negotiates IPsec security associations, and establishes IPsec keys. IKE uses RFCs 2408, 2409, 2410, and 2412, and additionally implements the draft-ietf-ipsec-ikev2-16.txt draft.

The term IPsec is sometimes used to describe the entire protocol of IPsec data services and IKE security protocols and is other times used to describe only the data services.

This chapter includes the following topics:

Information About IPsec Network Security

IP security (IPsec) protocol is a framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. It is developed by the Internet Engineering Task Force (IETF). IPsec provides security services at the IP layer, including protecting one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host. The overall IPsec implementation is the latest version of RFC 2401. Cisco NX-OS IPsec implements RFC 2402 through RFC 2410.

IPsec uses the Internet Key Exchange (IKE) protocol to handle protocol and algorithm negotiation and to generate the encryption and authentication keys used by IPsec. While IKE can be used with other protocols, its initial implementation is with the IPsec protocol. IKE provides authentication of the IPsec peers, negotiates IPsec security associations, and establishes IPsec keys. IKE uses RFCs 2408, 2409, 2410, and 2412, and additionally implements the draft-ietf-ipsec-ikev2-16.txt draft.

IPsec provides security for transmission of sensitive information over unprotected networks such as the Internet. IPsec acts at the network layer, protecting and authenticating IP packets between participating IPsec devices (peers).



Note IPsec is not supported by the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter.

IPsec provides security for transmission of sensitive information over unprotected networks such as the Internet. IPsec acts at the network layer, protecting and authenticating IP packets between participating IPsec devices (peers).

IPsec provides the following network security services. In general, the local security policy dictates the use of one or more of these services between two participating IPsec devices:

- Data confidentiality—The IPsec sender can encrypt packets before transmitting them across a network.
- Data integrity—The IPsec receiver can authenticate packets sent by the IPsec sender to ensure that the data has not been altered during transmission.
- Data origin authentication—The IPsec receiver can authenticate the source of the IPsec packets sent. This service is dependent upon the data integrity service.
- Anti-replay protection—The IPsec receiver can detect and reject replayed packets.



Note The term *data authentication* is generally used to mean data integrity and data origin authentication. Within this chapter it also includes anti-replay services, unless otherwise specified.

With IPsec, data can be transmitted across a public network without fear of observation, modification, or spoofing. This enables applications such as Virtual Private Networks (VPNs), including intranets, extranets, and remote user access.

IPsec as implemented in Cisco NX-OS software supports the Encapsulating Security Payload (ESP) protocol. This protocol encapsulates the data to be protected and provides data privacy services, optional data authentication, and optional anti-replay services.




Note The Encapsulating Security Payload (ESP) protocol is a header inserted into an existing TCP/IP packet, the size of which depends on the actual encryption and authentication algorithms negotiated. To avoid fragmentation, the encrypted packet fits into the interface maximum transmission unit (MTU). The path MTU calculation for TCP takes into account the addition of ESP headers, plus the outer IP header in tunnel mode, for encryption. The MDS switches allow 100 bytes for packet growth for IPsec encryption.



Note When using IPsec and IKE, each Gigabit Ethernet interface on the IPS module (either on 14+2 LC or 18+4 LC) must be configured in its own IP subnet. If there are multiple Gigabit Ethernet interfaces configured with IP address or network-mask in the same IP subnet, IKE packets may not be sent to the right peer and thus IPsec tunnel will not come up.

The following figure shows different IPsec scenarios.



This section includes the following topics:

About IKE

IKE automatically negotiates IPsec security associations and generates keys for all switches using the IPsec feature. Specifically, IKE provides these benefits:

- Allows you to refresh IPsec SAs.
- Allows IPsec to provide anti-replay services.
- Supports a manageable, scalable IPsec configuration.
- Allows dynamic authentication of peers.

IKE is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeSystem.

IPsec Compatibility

IPsec features are compatible with the following Cisco MDS 9000 Family hardware:

- Cisco 18/4-port Multi-Service Module (MSM-18/4) modules and MDS 9222i Module-1 modules.
- Cisco 14/2-port Multiprotocol Services (MPS-14/2) modules in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors
- Cisco MDS 9216i Switch with the 14/2-port multiprotocol capability in the integrated supervisor module. Refer to the *Cisco MDS 9200 Series Hardware Installation Guide* for more information on the Cisco MDS 9216i Switch.

- The IPsec feature is not supported on the management interface.

IPsec features are compatible with the following fabric setup:

- Two connected Cisco MDS 9200 Switches or Cisco MDS 9500 Directors running Cisco MDS SAN-OS Release 2.0(1b) or later, or Cisco NX-OS Release 4.1(1).
- A Cisco MDS 9200 Switches or Cisco MDS 9500 Directors running Cisco MDS SAN-OS Release 2.0(1b) or later, or Cisco NX-OS Release 4.1(1) connected to any IPsec compliant device.
- The following features are not supported in the Cisco NX-OS implementation of the IPsec feature:
 - Authentication Header (AH)
 - Transport mode
 - Security association bundling
 - Manually configuring security associations
 - Per host security association option in a crypto map
 - Security association idle timeout
 - Dynamic crypto maps

**Note**

Any reference to crypto maps in this document, only refers to static crypto maps.

IPsec and IKE Terminology

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

- Security association (SA)—An agreement between two participating peers on the entries required to encrypt and decrypt IP packets. Two SAs are required for each peer in each direction (inbound and outbound) to establish bidirectional communication between the peers. Sets of bidirectional SA records are stored in the SA database (SAD). IPsec uses IKE to negotiate and bring up SAs. Each SA record includes the following information:
 - Security parameter index (SPI)—A number which, together with a destination IP address and security protocol, uniquely identifies a particular SA. When using IKE to establish the SAs, the SPI for each SA is a pseudo-randomly derived number.
 - Peer—A switch or other device that participates in IPsec. For example, a Cisco MDS switch or other Cisco routers that support IPsec.
 - Transform—A list of operations done to provide data authentication and data confidentiality. For example, one transform is the ESP protocol with the HMAC-MD5 authentication algorithm.
 - Session key—The key used by the transform to provide security services.
 - Lifetime—A lifetime counter (in seconds and bytes) is maintained from the time the SA is created. When the time limit expires the SA is no longer operational and, if required, is automatically renegotiated (rekeyed).
 - Mode of operation—Two modes of operation are generally available for IPsec: tunnel mode and transport mode. The Cisco NX-OS implementation of IPsec only supports the tunnel mode. The IPsec tunnel mode encrypts and authenticates the IP packet, including its header. The gateways encrypt traffic on behalf of the hosts and subnets. The Cisco NX-OS implementation of IPsec does not support transport mode.



Note The term *tunnel mode* is different from the term *tunnel*, which is used to indicate a secure communication path between two peers, such as two switches connected by an FCIP link.

- Anti-replay—A security service where the receiver can reject old or duplicate packets to protect itself against replay attacks. IPsec provides this optional service by use of a sequence number combined with the use of data authentication.
- Data authentication—Data authentication can refer either to integrity alone or to both integrity and authentication (data origin authentication is dependent on data integrity).
 - Data integrity—Verifies that data has not been altered.
 - Data origin authentication—Verifies that the data was actually sent by the claimed sender.
- Data confidentiality—A security service where the protected data cannot be observed.
- Data flow—A grouping of traffic, identified by a combination of source address and mask or prefix, destination address mask or prefix length, IP next protocol field, and source and destination ports, where the protocol and port fields can have any of these values. Traffic matching a specific combination of these values is logically grouped together into a data flow. A data flow can represent a single TCP connection between two hosts, or it can represent traffic between two subnets. IPsec protection is applied to data flows.
- Perfect forward secrecy (PFS)—A cryptographic characteristic associated with a derived shared secret value. With PFS, if one key is compromised, previous and subsequent keys are not compromised, because subsequent keys are not derived from previous keys.
- Security Policy Database (SPD)—An ordered list of policies applied to traffic. A policy decides if a packet requires IPsec processing, if it should be allowed in clear text, or if it should be dropped.
 - The IPsec SPDs are derived from user configuration of crypto maps.
 - The IKE SPD is configured by the user.

Supported IPsec Transforms and Algorithms

The component technologies implemented for IPsec include the following transforms:

- Advanced Encrypted Standard (AES) is an encryption algorithm. It implements either 128 or 256 bits using Cipher Block Chaining (CBC) or counter mode.
- Data Encryption Standard (DES) is used to encrypt packet data and implements the mandatory 56-bit DES-CBC. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPsec packet.
- Triple DES (3DES) is a stronger form of DES with 168-bit encryption keys that allow sensitive information to be transmitted over untrusted networks.



Note Cisco NX-OS images with strong encryption are subject to United States government export controls, and have a limited distribution. Images to be installed outside the United States require an export license. Customer orders might be denied or subject to delay due to United States government regulations. Contact your sales representative or distributor for more information, or send e-mail to export@cisco.com.

- Message Digest 5 (MD5) is a hash algorithm with the HMAC variant. HMAC is a keyed hash variant used to authenticate data.

- Secure Hash Algorithm (SHA-1) is a hash algorithm with the Hash Message Authentication Code (HMAC) variant.
- AES-XCBC-MAC is a Message Authentication Code (MAC) using the AES algorithm.

Supported IKE Transforms and Algorithms

The component technologies implemented for IKE include the following transforms:

- Diffie-Hellman (DH) is a public-key cryptography protocol that allows two parties to establish a shared secret over an unsecure communications channel. Diffie-Hellman is used within IKE to establish session keys. Group 1 (768-bit), Group 2 (1024-bit), and Group 5 (1536-bit) are supported.
- Advanced Encrypted Standard (AES) is an encryption algorithm. It implements either 128 bits using Cipher Block Chaining (CBC) or counter mode.
- Data Encryption Standard (DES) is used to encrypt packet data and implements the mandatory 56-bit DES-CBC. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPsec packet.
- Triple DES (3DES) is a stronger form of DES with 168-bit encryption keys that allow sensitive information to be transmitted over untrusted networks.



Note

Cisco NX-OS images with strong encryption are subject to United States government export controls, and have a limited distribution. Images to be installed outside the United States require an export license. Customer orders might be denied or subject to delay due to United States government regulations. Contact your sales representative or distributor for more information, or send e-mail to export@cisco.com.

- Message Digest 5 (MD5) is a hash algorithm with the HMAC variant. HMAC is a keyed hash variant used to authenticate data.
- Secure Hash Algorithm (SHA-1) is a hash algorithm with the Hash Message Authentication Code (HMAC) variant.
- The switch authentication algorithm uses the preshared keys based on the IP address


About IPsec Digital Certificate Support

This section describes the advantages of using certificate authorities (CAs) and digital certificates for authentication.

Implementing IPsec Without CAs and Digital Certificates


Without a CA and digital certificates, enabling IPsec services (such as encryption) between two Cisco MDS switches requires that each switch has the key of the other switch (such as an RSA public key or a shared key). You must manually specify either the RSA public keys or preshared keys on each switch in the fabric using IPsec services. Also, each new device added to the fabric will require manual configuration of the other switches in the fabric to support secure communication. Each switch uses the key of the other switch to authenticate the identity of the other switch; this authentication always occurs when IPsec traffic is exchanged between the two switches.

If you have multiple Cisco MDS switches in a mesh topology and want to exchange IPsec traffic passing among all of those switches, you must first configure shared keys or RSA public keys among all of those switches.



Every time a new switch is added to the IPsec network, you must configure keys between the new switch and each of the existing switches. In the following figure, four additional two-part key configurations are required to add a single encrypting switch to the network.


Consequently, the more devices that require IPsec services, the more involved the key administration becomes. This approach does not scale well for larger, more complex encrypting networks.



Implementing IPsec with CAs and Digital Certificates

With CA and digital certificates, you do not have to configure keys between all the encrypting switches. Instead, you individually enroll each participating switch with the CA, requesting a certificate for the switch. When this has been accomplished, each participating switch can dynamically authenticate all the other participating switches. When two devices want to communicate, they exchange certificates and digitally sign data to authenticate each other. When a new device is added to the network, you simply enroll that device with a CA, and none of the other devices needs modification. When the new device attempts an IPsec connection, certificates are automatically exchanged and the device can be authenticated.

The following figure shows the process of dynamically authenticating the devices.



To add a new IPsec switch to the network, you need only configure that new switch to request a certificate from the CA, instead of making multiple key configurations with all the other existing IPsec switches.

How CA Certificates Are Used by IPsec Devices

When two IPsec switches want to exchange IPsec-protected traffic passing between them, they must first authenticate each other—otherwise, IPsec protection cannot occur. The authentication is done with IKE.

IKE can use two methods to authenticate the switches, using preshared keys without a CA and using RSA key-pairs with a CA. Both methods require that keys must be preconfigured between the two switches.

Without a CA, a switch authenticates itself to the remote switch using either RSA-encrypted preshared keys.

With a CA, a switch authenticates itself to the remote switch by sending a certificate to the remote switch and performing some public key cryptography. Each switch must send its own unique certificate that was issued and validated by the CA. This process works because the certificate of each switch encapsulates the public key of the switch, each certificate is authenticated by the CA, and all participating switches recognize the CA as an authenticating authority. This scheme is called IKE with an RSA signature.

Your switch can continue sending its own certificate for multiple IPsec sessions, and to multiple IPsec peers until the certificate expires. When the certificate expires, the switch administrator must obtain a new one from the CA.

CAs can also revoke certificates for devices that will no longer participate in IPsec. Revoked certificates are not recognized as valid by other IPsec devices. Revoked certificates are listed in a certificate revocation list (CRL), which each peer may check before accepting a certificate from another peer.

Certificate support for IKE has the following considerations:

- The switch FQDN (host name and domain name) must be configured before installing certificates for IKE.
- Only those certificates that are configured for IKE or general usage are used by IKE.
- The first IKE or general usage certificate configured on the switch is used as the default certificate by IKE.
- The default certificate is for all IKE peers unless the peer specifies another certificate.
- If the peer asks for a certificate which is signed by a CA that it trusts, then IKE uses that certificate, if it exists on the switch, even if it is not the default certificate.
- If the default certificate is deleted, the next IKE or general usage certificate, if any exists, is used by IKE as the default certificate.

- Certificate chaining is not supported by IKE.
- IKE only sends the identity certificate, not the entire CA chain. For the certificate to be verified on the peer, the same CA chain must also exist there.

About IKE Initialization

The IKE feature must first be enabled and configured so the IPsec feature can establish data flow with the required peer. DCNM-SAN initializes IKE when you first configure it.

You cannot disable IKE if IPsec is enabled. If you disable the IKE feature, the IKE configuration is cleared from the running configuration.

About the IKE Domain

You must apply the IKE configuration to an IPsec domain to allow traffic to reach the supervisor module in the local switch. DCNM-SAN sets the IPsec domain automatically when you configure IKE.

About IKE Tunnels

An IKE tunnel is a secure IKE session between two endpoints. IKE creates this tunnel to protect IKE messages used in IPsec SA negotiations.

Two versions of IKE are used in the Cisco NX-OS implementation.

- IKE version 1 (IKEv1) is implemented using RFC 2407, 2408, 2409, and 2412.
- IKE version 2 (IKEv2) is a simplified and more efficient version and does not interoperate with IKEv1. IKEv2 is implemented using the draft-ietf-ipsec-ikev2-16.txt draft.

About IKE Policy Negotiation

To protect IKE negotiations, each IKE negotiation begins with a common (shared) IKE policy. An IKE policy defines a combination of security parameters to be used during the IKE negotiation. By default, no IKE policy is configured. You must create IKE policies at each peer. This policy states which security parameters will be used to protect subsequent IKE negotiations and mandates how peers are authenticated. You can create multiple, prioritized policies at each peer to ensure that at least one policy will match a remote peer's policy.

You can configure the policy based on the encryption algorithm (DES, 3DES, or AES), the hash algorithm (SHA or MD5), and the DH group (1, 2, or 5). Each policy can contain a different combination of parameter values. A unique priority number identifies the configured policy. This number ranges from 1 (highest priority) to 255 (lowest priority). You can create multiple policies in a switch. If you need to connect to a remote peer, you must ascertain that at least one policy in the local switch contains the identical parameter values configured in the remote peer. If several policies have identical parameter configurations, the policy with the lowest number is selected.

[Table 86: IKE Transform Configuration Parameters](#), on page 715 provides a list of allowed transform combinations.

Table 86: IKE Transform Configuration Parameters

Parameter	Accepted Values	Keyword	Default Value
encryption algorithm	56-bit DES-CBC	des	3des
	168-bit DES	3des	
	128-bit AES	aes	

Parameter	Accepted Values	Keyword	Default Value
hash algorithm	SHA-1 (HMAC variant) MD5 (HMAC variant)	sha md5	sha
authentication method	Preshared keys	Not configurable	Preshared keys
DH group identifier	768-bit DH 1024-bit DH 1536-bit DH	1 2 5	1

The following table lists the supported and verified settings for IPsec and IKE encryption authentication algorithms on the Microsoft Windows and Linux platforms:

Platform	IKE	IPsec
Microsoft iSCSI initiator, Microsoft IPsec implementation on Microsoft Windows 2000 platform	3DES, SHA-1 or MD5, DH group 2	3DES, SHA-1
Cisco iSCSI initiator, Free Swan IPsec implementation on Linux platform	3DES, MD5, DH group 1	3DES, MD5

**Note**

When you configure the hash algorithm, the corresponding HMAC version is used as the authentication algorithm.

When the IKE negotiation begins, IKE looks for an IKE policy that is the same on both peers. The peer that initiates the negotiation will send all its policies to the remote peer, and the remote peer will try to find a match. The remote peer looks for a match by comparing its own highest priority policy against the other peer's received policies. The remote peer checks each of its policies in order of its priority (highest priority first) until a match is found.

A match is found when the two peers have the same encryption, hash algorithm, authentication algorithm, and DH group values. If a match is found, IKE completes the security negotiation and the IPsec SAs are created.

If an acceptable match is not found, IKE refuses negotiation and the IPsec data flows will not be established.

Optional IKE Parameter Configuration

You can optionally configure the following parameters for the IKE feature:

- The lifetime association within each policy—The lifetime ranges from 600 to 86,400 seconds. The default is 86,400 seconds (equals one day). The lifetime association within each policy is configured when you are creating an IKE policy. See the [Configuring an IKE Policy, on page 729](#).
- The keepalive time for each peer if you use IKEv2—The keepalive ranges from 120 to 86,400 seconds. The default is 3,600 seconds (equals one hour).
- The initiator version for each peer—IKE v1 or IKE v2 (default). Your choice of initiator version does not affect interoperability when the remote device initiates the negotiation. Configure this option if the

peer device supports IKEv1 and you can play the initiator role for IKE with the specified device. Use the following considerations when configuring the initiator version with FCIP tunnels:

- If the switches on both sides of an FCIP tunnel are running MDS SAN-OS Release 3.0(1) or later, or Cisco NX-OS 4.1(1) you must configure initiator version IKEv1 on both sides of an FCIP tunnel to use only IKEv1. If one side of an FCIP tunnel is using IKEv1 and the other side is using IKEv2, the FCIP tunnel uses IKEv2.
- If the switch on one side of an FCIP tunnel is running MDS SAN-OS Release 3.0(1) or later, or Cisco NX-OS 4.1(1b) and the switch on the other side of the FCIP tunnel is running MDS SAN-OS Release 2.x, configuring IKEv1 on either side (or both) results in the FCIP tunnel using IKEv1.



Note Only IKE v1 is supported to build IPsec between 2.x and 3.x MDS switches.



Caution You may need to configure the initiator version even when the switch does not behave as an IKE initiator under normal circumstances. Always using this option guarantees a faster recovery of traffic flows in case of failures.



Tip The keepalive time only applies to IKEv2 peers and not to all peers.



Note When IPsec implementations in the host prefer to initiate the IPsec rekey, be sure to configure the IPsec lifetime value in the Cisco MDS switch to be higher than the lifetime value in the host.

About Crypto IPv4-ACLs

IP access control lists (IPv4-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IPv4 IP-ACLs restrict IP-related traffic based on the configured IP filters. See the *Configuring IPv4 and IPv6 Access Control Lists* chapter for details on creating and defining IPv4-ACLs.

In the context of crypto maps, IPv4-ACLs are different from regular IPv4-ACLs. Regular IPv4-ACLs determine what traffic to forward or block at an interface. For example, IPv4-ACLs can be created to protect all IP traffic between subnet A and subnet Y or Telnet traffic between host A and host B.

Crypto IPv4-ACLs are used to define which IP traffic requires crypto protection and which traffic does not.

Crypto IPv4-ACLs associated with IPsec crypto map entries have four primary functions:

- Select outbound traffic to be protected by IPsec (permit = protect).
- Indicate the data flow to be protected by the new SAs (specified by a single permit entry) when initiating negotiations for IPsec SAs.
- Process inbound traffic to filter out and discard traffic that should have been protected by IPsec.
- Determine whether or not to accept requests for IPsec SAs on behalf of the requested data flows when processing IKE negotiation from the IPsec peer.

**Tip**

If you want some traffic to receive one type of IPsec protection (for example, encryption only) and other traffic to receive a different type of IPsec protection (for example, both authentication and encryption), create two IPv4-ACLs. Use both IPv4-ACLs in different crypto maps to specify different IPsec policies.

**Note**

IPsec does not support IPv6-ACLs.

Mirror Image Crypto IPv4-ACLs

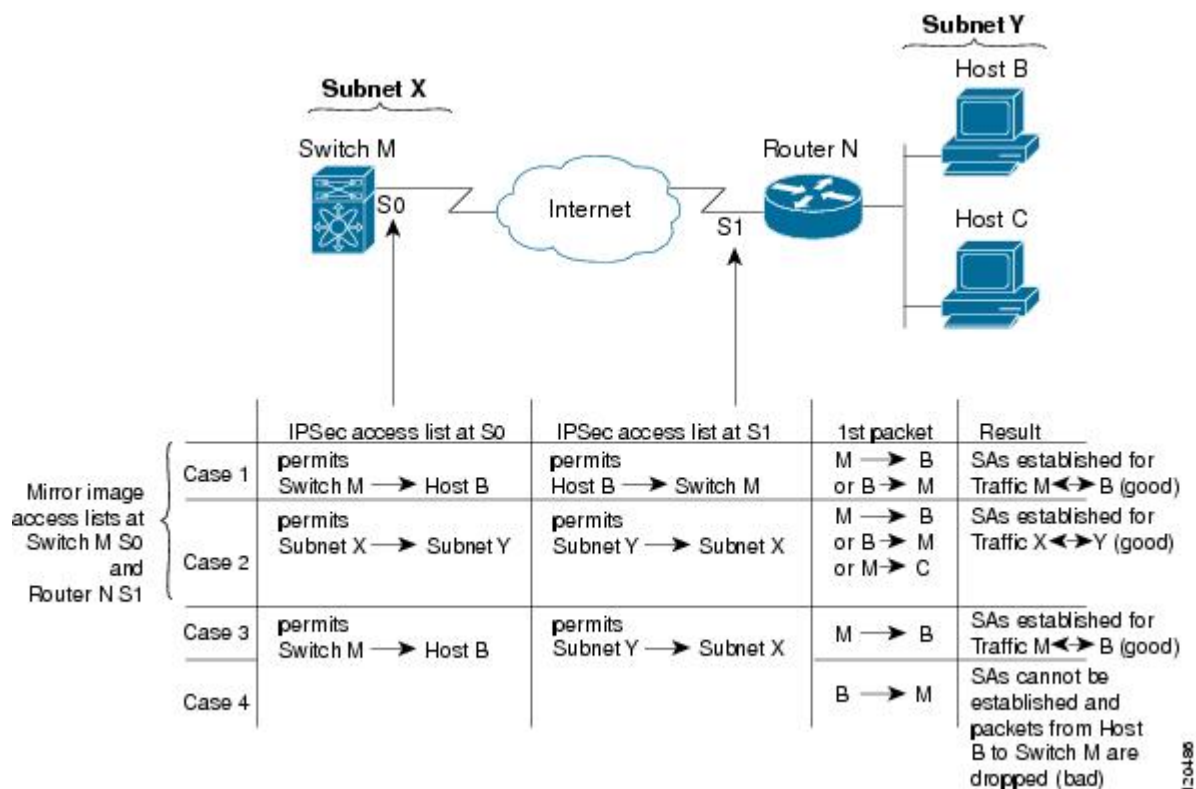
For every crypto IPv4-ACL specified for a crypto map entry defined at the local peer, define a mirror image crypto IPv4-ACL at the remote peer. This configuration ensures that IPsec traffic applied locally can be processed correctly at the remote peer.

**Tip**

The crypto map entries themselves must also support common transforms and must refer to the other system as a peer.

Figure 95: IPsec Processing of Mirror Image Configuration, on page 718 shows some sample scenarios with and without mirror image IPv4-ACLs.

Figure 95: IPsec Processing of Mirror Image Configuration



As [Figure 95: IPsec Processing of Mirror Image Configuration, on page 718](#) indicates, IPsec SAs can be established as expected whenever the two peers' crypto IPv4-ACLs are mirror images of each other. However, an IPsec SA can be established only some of the time when the IPv4-ACLs are not mirror images of each other. This can happen in the case when an entry in one peer's IPv4-ACL is a subset of an entry in the other peer's IPv4-ACL, such as shown in cases 3 and 4 of [Figure 95: IPsec Processing of Mirror Image Configuration, on page 718](#). IPsec SA establishment is critical to IPsec. Without SAs, IPsec does not work, causing any packets matching the crypto IPv4-ACL criteria to be silently dropped instead of being forwarded with IPsec security.

In case 4, an SA cannot be established because SAs are always requested according to the crypto IPv4-ACLs at the initiating packet's end. In case 4, router N requests that all traffic between subnet X and subnet Y be protected, but this is a superset of the specific flows permitted by the crypto IPv4-ACL at switch M so the request is not permitted. Case 3 works because switch M's request is a subset of the specific flows permitted by the crypto IPv4-ACL at router N.

Because of the complexities introduced when crypto IPv4-ACLs are not configured as mirror images at peer IPsec devices, we strongly encourage you to use mirror image crypto IPv4-ACLs.

The any Keyword in Crypto IPv4-ACLs



Tip We recommend that you configure mirror image crypto IPv4-ACLs for use by IPsec and that you avoid using the **any** option.

The **any** keyword in a permit statement is discouraged when you have multicast traffic flowing through the IPsec interface. This configuration can cause multicast traffic to fail.

The **permit any** statement causes all outbound traffic to be protected (and all protected traffic sent to the peer specified in the corresponding crypto map entry) and requires protection for all inbound traffic. Then, all inbound packets that lack IPsec protection are silently dropped, including packets for routing protocols, NTP, echo, echo response, and so forth.

You need to be sure you define which packets to protect. If you must use **any** in a permit statement, you must preface that statement with a series of deny statements to filter out any traffic (that would otherwise fall within that permit statement) that you do not want to be protected.

About Transform Sets in IPsec

A transform set represents a certain combination of security protocols and algorithms. During the IPsec security association negotiation, the peers agree to use a particular transform set for protecting a particular data flow.

You can specify multiple transform sets, and then specify one or more of these transform sets in a crypto map entry. The transform set defined in the crypto map entry is used in the IPsec security association negotiation to protect the data flows specified by that crypto map entry's access list.

During IPsec security association negotiations with IKE, the peers search for a transform set that is the same at both peers. When such a transform set is found, it is selected and applied to the protected traffic as part of both peers' IPsec security associations.



Tip If you change a transform set definition, the change is only applied to crypto map entries that reference the transform set. The change is not applied to existing security associations, but used in subsequent negotiations to establish new security associations. If you want the new settings to take effect sooner, you can clear all or part of the security association database.



Note When you enable IPsec, the Cisco NX-OS software automatically creates a default transform set (ipsec_default_tranform_set) using AES-128 encryption and SHA-1 authentication algorithms.

[Table 87: IPsec Transform Configuration Parameters, on page 720](#) provides a list of allowed transform combinations for IPsec.

Table 87: IPsec Transform Configuration Parameters

Parameter	Accepted Values	Keyword
encryption algorithm	56-bit DES-CBC 168-bit DES 128-bit AES-CBC 128-bit AES-CTR ³⁴ 256-bit AES-CBC 256-bit AES-CTR 1	esp-des esp-3des esp-aes 128 esp-aes 128 ctr esp-aes 256 esp-aes 256 ctr
hash/authentication algorithm 1 (optional)	SHA-1 (HMAC variant) MD5 (HMAC variant) AES-XCBC-MAC	esp-sha1-hmac esp-md5-hmac esp-aes-xcbc-mac

³⁴ If you configure the AES counter (CTR) mode, you must also configure the authentication algorithm.

The following table lists the supported and verified settings for IPsec and IKE encryption authentication algorithms on the Microsoft Windows and Linux platforms:

Platform	IKE	IPsec
Microsoft iSCSI initiator, Microsoft IPsec implementation on Microsoft Windows 2000 platform	3DES, SHA-1 or MD5, DH group 2	3DES, SHA-1
Cisco iSCSI initiator, Free Swan IPsec implementation on Linux platform	3DES, MD5, DH group 1	3DES, MD5

About Crypto Map Entries

Once you have created the crypto IPv4-ACLs and transform sets, you can create crypto map entries that combine the various parts of the IPsec SA, including the following:

- The traffic to be protected by IPsec (per the crypto IPv4-ACL). A crypto map set can contain multiple entries, each with a different IPv4-ACL.
- The granularity of the flow to be protected by a set of SAs.
- The IPsec-protected traffic destination (who the remote IPsec peer is).
- The local address to be used for the IPsec traffic (applying to an interface).
- The IPsec security to be applied to this traffic (selecting from a list of one or more transform sets).
- Other parameters to define an IPsec SA.

Crypto map entries with the same crypto map name (but different map sequence numbers) are grouped into a crypto map set.

When you apply a crypto map set to an interface, the following events occur:

- A security policy database (SPD) is created for that interface.
- All IP traffic passing through the interface is evaluated against the SPD.

If a crypto map entry sees outbound IP traffic that requires protection, an SA is negotiated with the remote peer according to the parameters included in the crypto map entry.

The policy derived from the crypto map entries is used during the negotiation of SAs. If the local switch initiates the negotiation, it will use the policy specified in the crypto map entries to create the offer to be sent to the specified IPsec peer. If the IPsec peer initiates the negotiation, the local switch checks the policy from the crypto map entries and decides whether to accept or reject the peer's request (offer).

For IPsec to succeed between two IPsec peers, both peers' crypto map entries must contain compatible configuration statements.

SA Establishment Between Peers

When two peers try to establish an SA, they must each have at least one crypto map entry that is compatible with one of the other peer's crypto map entries.

For two crypto map entries to be compatible, they must at least meet the following criteria:

- The crypto map entries must contain compatible crypto IPv4-ACLs (for example, mirror image IPv4-ACLs). If the responding peer entry is in the local crypto, the IPv4-ACL must be permitted by the peer's crypto IPv4-ACL.
- The crypto map entries must each identify the other peer or must have auto peer configured.
- If you create more than one crypto map entry for a given interface, use the seq-num of each map entry to rank the map entries: the lower the seq-num, the higher the priority. At the interface that has the crypto map set, traffic is evaluated against higher priority map entries first.
- The crypto map entries must have at least one transform set in common, where IKE negotiations are carried out and SAs are established. During the IPsec SA negotiation, the peers agree to use a particular transform set when protecting a particular data flow.

When a packet matches a permit entry in a particular IPv4-ACL, the corresponding crypto map entry is tagged, and the connections are established.

About SA Lifetime Negotiation

You can override the global lifetime values (size and time) by configuring an SA-specific lifetime value.

To specify SA lifetime negotiation values, you can optionally configure the lifetime value for a specified crypto map. If you do, this value overrides the globally set values. If you do not specify the crypto map specific lifetime, the global value (or global default) is used.


See the [Global Lifetime Values, on page 723](#) for more information on global lifetime values.

About the AutoPeer Option

Setting the peer address as **auto-peer** in the crypto map indicates that the destination endpoint of the traffic should be used as the peer address for the SA. Using the same crypto map, a unique SA can be set up at each of the endpoints in the subnet specified by the crypto map's IPv4-ACL entry. Auto-peer simplifies configuration when traffic endpoints are IPsec capable. It is particularly useful for iSCSI, where the iSCSI hosts in the same subnet do not require separate configuration.

Below figure shows a scenario where the auto-peer option can simplify configuration. Using the auto-peer option, only one crypto map entry is needed for all the hosts from subnet X to set up SAs with the switch. Each host will set up its own SA, but will share the crypto map entry. Without the auto-peer option, each host needs one crypto map entry.

Figure 96: iSCSI with End-to-End IPsec Using the auto-peer Option



About Perfect Forward Secrecy

To specify SA lifetime negotiation values, you can also optionally configure the perfect forward secrecy (PFS) value in the crypto map.

The PFS feature is disabled by default. If you set the PFS group, you can set one of the DH groups: 1, 2, 5, or 14. If you do not specify a DH group, the software uses group 1 by default.

About Crypto Map Set Interface Application

You need to apply a crypto map set to each interface through which IPsec traffic will flow. Applying the crypto map set to an interface instructs the switch to evaluate all the interface's traffic against the crypto map set and to use the specified policy during connection or SA negotiation on behalf of the traffic to be protected by crypto.

You can apply only one crypto map set to an interface. You can apply the same crypto map to multiple interfaces. However, you cannot apply more than one crypto map set to each interface.

IPsec Maintenance

Certain configuration changes will only take effect when negotiating subsequent security associations. If you want the new settings to take immediate effect, you must clear the existing security associations so that they will be reestablished with the changed configuration. If the switch is actively processing IPsec traffic, it is

desirable to clear only the portion of the security association database that would be affected by the configuration changes (that is, clear only the security associations established by a given crypto map set). Clearing the full security association database should be reserved for large-scale changes, or when the router is processing very little other IPsec traffic.



Tip You can obtain the SA index from the output of the **show crypto sa domain interface gigabitethernet slot/port** command.

Use the following command to clear part of the SA database.

```
switch# clear crypto sa domain ipsec interface gigabitethernet 2/1 inbound sa-index 1
```

Global Lifetime Values

If you have not configured a lifetime in the crypto map entry, the global lifetime values are used when negotiating new IPsec SAs.

You can configure two lifetimes: timed or traffic-volume. An SA expires after the first of these lifetimes is reached. The default lifetimes are 3,600 seconds (one hour) and 450 GB.

If you change a global lifetime, the new lifetime value will not be applied to currently existing SAs, but will be used in the negotiation of subsequently established SAs. If you wish to use the new values immediately, you can clear all or part of the SA database.

Assuming that the particular crypto map entry does not have lifetime values configured, when the switch requests new SAs it will specify its global lifetime values in the request to the peer; it will use this value as the lifetime of the new SAs. When the switch receives a negotiation request from the peer, it uses the value determined by the IKE version in use:

- If you use IKEv1 to set up IPsec SAs, the SA lifetime values are chosen to be the smaller of the two proposals. The same values are programmed on both the ends of the tunnel.
- If you use IKEv2 to set up IPsec SAs, the SAs on each end have their own set up of lifetime values and thus the SAs on both sides expire independently.

The SA (and corresponding keys) will expire according to whichever comes sooner, either after the specified amount of time (in seconds) has passed or after the specified amount of traffic (in bytes) has passed.

A new SA is negotiated before the lifetime threshold of the existing SA is reached to ensure that negotiation completes before the existing SA expires.

The new SA is negotiated when one of the following thresholds is reached (whichever comes first):

- 30 seconds before the lifetime expires or
- Approximately 10% of the lifetime in bytes remain

If no traffic has passed through when the lifetime expires, a new SA is not negotiated. Instead, a new SA will be negotiated only when IPsec sees another packet that should be protected.

Prerequisites for IPsec

To use the IPsec feature, you need to perform the following tasks:

- Obtain the ENTERPRISE_PKG license (see the Cisco MDS 9000 Family NX-OS Licensing Guide).
- Configure IKE as described in the [About IKE Initialization, on page 715](#).

Guidelines and Limitations

The following are the guidelines and limitations for IPsec network security:

Crypto IPv4-ACL Guidelines

Follow these guidelines when configuring IPv4-ACLs for the IPsec feature:

- The Cisco NX-OS software only allows name-based IPv4-ACLs.
- When an IPv4-ACL is applied to a crypto map, the following options apply:
 - Permit—Applies the IPsec feature to the traffic.
 - Deny—Allows clear text (default).



Note

IKE traffic (UDP port 500) is implicitly transmitted in clear text.

- The IPsec feature only considers the source and destination IPv4 addresses and subnet masks, protocol, and single port number. There is no support for IPv6 in IPsec.



Note

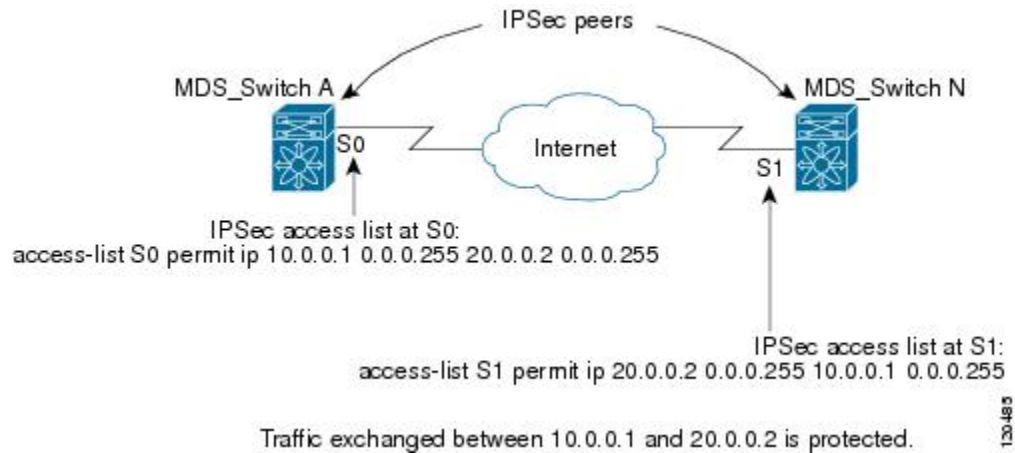
The IPsec feature does not support port number ranges and ignores higher port number field, if specified.

- The permit option causes all IP traffic that matches the specified conditions to be protected by crypto, using the policy described by the corresponding crypto map entry.
- The deny option prevents traffic from being protected by crypto. The first deny statement causes the traffic to be in clear text.
- The crypto IPv4-ACL you define is applied to an interface after you define the corresponding crypto map entry and apply the crypto map set to the interface.
- Different IPv4-ACLs must be used in different entries of the same crypto map set.
- Inbound and outbound traffic is evaluated against the same outbound IPv4-ACL. Therefore, the IPv4-ACL's criteria is applied in the forward direction to traffic exiting your switch, and the reverse direction to traffic entering your switch.
- Each IPv4-ACL filter assigned to the crypto map entry is equivalent to one security policy entry. The IPsec feature supports up to 120 security policy entries for each MPS-14/2 module and Cisco MDS 9216i Switch.
- IPsec protection (see [Figure 97: IPsec Processing of Crypto IPv4-ACLs](#), on page 725) is applied to traffic between switch interface S0 (IPv4 address 10.0.0.1) and switch interface S1 (IPv4 address 20.0.0.2) as the data exits switch A's S0 interface enroute to switch interface S1. For traffic from 10.0.0.1 to 20.0.0.2, the IPv4-ACL entry on switch A is evaluated as follows:
 - source = IPv4 address 10.0.0.1
 - dest = IPv4 address 20.0.0.2

For traffic from 20.0.0.2 to 10.0.0.1, that same IPv4-ACL entry on switch A is evaluated as follows:

- source = IPv4 address 20.0.0.2
- dest = IPv4 address 10.0.0.1

Figure 97: IPsec Processing of Crypto IPv4-ACLs



- If you configure multiple statements for a given crypto IPv4-ACL that is used for IPsec, the first permit statement that is matched is used to determine the scope of the IPsec SA. Later, if traffic matches a different permit statement of the crypto IPv4-ACL, a new, separate IPsec SA is negotiated to protect traffic matching the newly matched IPv4-ACL statement.
- Unprotected inbound traffic that matches a permit entry in the crypto IPv4-ACL for a crypto map entry flagged as IPsec is dropped, because this traffic was expected to be protected by IPsec.
- You can use the **show ip access-lists** command to view all IP-ACLs. The IP-ACLs used for traffic filtering purposes are also used for crypto.
- For IPsec to interoperate effectively with Microsoft iSCSI initiators, specify the TCP protocol and the local iSCSI TCP port number (default 3260) in the IPv4-ACL. This configuration ensures the speedy recovery of encrypted iSCSI sessions following disruptions such as Gigabit Ethernet interfaces shutdowns, VRRP switchovers, and port failures.
- The following example of a IPv4-ACL entry shows that the MDS switch IPv4 address is 10.10.10.50 and remote Microsoft host running encrypted iSCSI sessions is 10.10.10.16:

```
switch(config)# ip access-list aclmsiscsi2 permit tcp 10.10.10.50 0.0.0.0 range port 3260
3260 10.10.10.16 0.0.0.0
```

Crypto Map Configuration Guidelines

When configuring crypto map entries, follow these guidelines:

- The sequence number for each crypto map decides the order in which the policies are applied. A lower sequence number is assigned a higher priority.
- Only one IPv4-ACL is allowed for each crypto map entry (the IPv4-ACL itself can have multiple permit or deny entries).
- When the tunnel endpoint is the same as the destination address, you can use the auto-peer option to dynamically configure the peer.
- For IPsec to interoperate effectively with Microsoft iSCSI initiators, specify the TCP protocol and the local iSCSI TCP port number (default 3260) in the IPv4-ACL. This configuration ensures the speedy recovery of encrypted iSCSI sessions following disruptions such as Gigabit Ethernet interfaces shutdowns, VRRP switchovers, and port failures.

Default Settings

[Table 88: Default IKE Parameters](#) , on page 726 lists the default settings for IKE parameters.

Table 88: Default IKE Parameters

Parameters	Default
IKE	Disabled.
IKE version	IKE version 2.
IKE encryption algorithm	3DES.
IKE hash algorithm	SHA.
IKE authentication method	Not configurable (uses preshared Preshared keys).
IKE DH group identifier	Group 1.
IKE lifetime association	86,400 00 seconds (equals 24 hours).
IKE keepalive time for each peer (v2)	3,600 seconds (equals 1 hour).

[Table 89: Default IPsec Parameters](#) , on page 726 lists the default settings for IPsec parameters.

Table 89: Default IPsec Parameters

Parameters	Default
IPsec	Disabled.
Applying IPsec to the traffic.	Deny—allowing clear text.
IPsec PFS	Disabled.
IPsec global lifetime (traffic-volume)	450 Gigabytes.
IPsec global lifetime (time)	3,600 seconds (one hour).

Enabling IPsec Using FCIP Wizard

DCNM-SAN simplifies the configuration of IPsec and IKE by enabling and configuring these features as part of the FCIP configuration using the FCIP Wizard.

To enable IPsec using the FCIP Wizard, follow these steps:

Procedure

Step 1

Click the FCIP Wizard icon in the toolbar.

Figure 98: FCIP Wizard



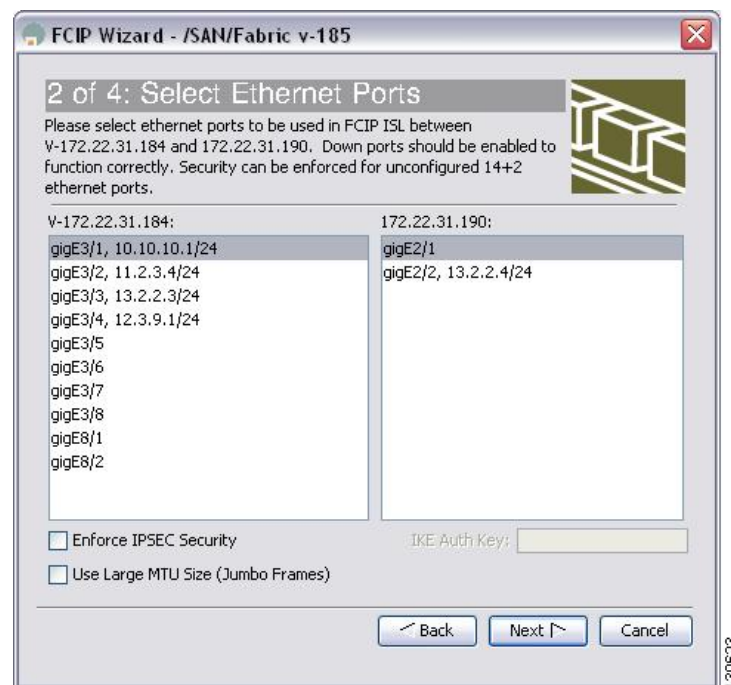
Step 2 Choose the switches that act as end points for the FCIP link and click Next.

Note These switches must have MPS-14/2 modules installed to configure IPsec on this FCIP link.

Step 3 Choose the Gigabit Ethernet ports on each MPS-14/2 module that will form the FCIP link.

Step 4 Check the Enforce IPSEC Security check box and set IKE Auth Key (see [Figure 99: Enabling IPsec on an FCIP Link](#), on page 727).

Figure 99: Enabling IPsec on an FCIP Link



Step 5 Click Next. In the Specify Tunnel Properties dialog box, you see the TCP connection characteristics.

Step 6 Set the minimum and maximum bandwidth settings and round-trip time for the TCP connections on this FCIP link. Click the Measure button to measure the round-trip time between the Gigabit Ethernet endpoints.

Step 7 Check the Enable Write Acceleration check box to enable FCIP write acceleration on this FCIP link.

Step 8 Check the Enable Optimum Compression check box to enable IP compression on this FCIP link.

Step 9 Click Next to configure the FCIP tunnel parameters.

Step 10 Set the Port VSAN for nontrunk/auto and allowed VSAN list for the trunk tunnel. Choose a Trunk Mode for this FCIP link. See the IP Services Configuration Guide, Cisco DCNM for SAN.

Step 11 Click Finish to create this FCIP link or click Cancel to exit the FCIP Wizard without creating an FCIP link.

Verifying IPsec and IKE

To verify that IPsec and IKE are enabled, follow these steps:

Procedure

-
- | | |
|---------------|-------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Expand Switches > Security and then select IPSEC in the Physical Attributes pane. |
| Step 2 | The Control tab is the default. Verify that the switches you want to modify for IPsec are enabled in the Status column. |
| Step 3 | Expand Switches > Security, and then select IKE in the Physical Attributes pane. |
| Step 4 | The Control tab is the default. Verify that the switches you want to modify for IKE are enabled in the Status column. |
-

Configuring IPsec and IKE Manually

This section describes how to manually configure IPsec and IKE.

If you are not using the FCIP Wizard, see [Enabling IPsec Using FCIP Wizard, on page 726](#).

IPsec provides secure data flows between participating peers. Multiple IPsec data flows can exist between two peers to secure different data flows, with each tunnel using a separate set of SAs.

Before you begin

After you have completed IKE configuration, configure IPsec.

To configure IPsec in each participating IPsec peer, follow these steps:

Procedure

-
- | | |
|---------------|-----------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Identify the peers for the traffic to which secure tunnels should be established. |
| Step 2 | Configure the transform set with the required protocols and algorithms. |
| Step 3 | Create the crypto map and apply access control lists (IPv4-ACLs), transform sets, peers, and lifetime values as applicable. |
| Step 4 | Apply the crypto map to the required interface. |
-

What to do next

This section includes the following topics:

Using IPsec

To use the IPsec feature, follow these steps:

Procedure

-
- Step 1** Obtain the ENTERPRISE_PKG license to enable IPsec for iSCSI and to enable IPsec for FCIP. See the Cisco MDS 9000 Family NX-OS Licensing Guide for more information.
- Step 2** Configure IKE as described in the [Configuring IPsec and IKE Manually, on page 728](#).
- Note** The IPsec feature inserts new headers in existing packets (see the Cisco MDS 9000 Family NX-OS IP Services Configuration Guide and Cisco DCNM for SAN for more information).
-

Configuring an IKE Policy

To configure the IKE policy negotiation parameters, follow these steps:

Procedure

-
- Step 1** Expand Switches > Security, and then select IKE.
- Step 2** Click the Policies tab.
- You see the existing IKE policies in the Information pane.
- Step 3** Click Create Row to create an IKE policy.
- Step 4** Enter the Priority for this switch. You can enter a value from one through 255, one being the highest.
- Step 5** Select appropriate values for the encryption, hash, authentication, and DHGroup fields.
- Step 6** Enter the lifetime for the policy. You can enter a lifetime from 600 to 86400 seconds.
- Step 7** Click Create to create this policy, or click Close to discard any unsaved changes.
-

Configuring the Keepalive Time for a Peer

To configure the keepalive time for each peer, follow these steps:

Procedure

-
- Step 1** Expand Switches > Security, and then select IKE.
- Step 2** Click the Global tab.
- Step 3** Enter a value (in seconds) in the KeepAliveInterval (sec). The keepalive interval in seconds is used by the IKE entity on the managed device with all the peers for the DOI corresponding to this conceptual row.
- Step 4** Click **Apply Changes** to save your changes.
-

Configuring the Initiator Version

To configure the initiator version, follow these steps:

Procedure

-
- Step 1** Expand **Switches > Security**, and then select **IKE**.
- Step 2** Click the **Initiator Version** tab.
- You see the existing initiator versions for the peers in the Information pane.
- Step 3** Click **Create Row** to create an initiator version.
- Step 4** Select the **Switches** for the remote peer for which this IKE protocol initiator is configured.
- Step 5** Enter the IP address of the remote peer.
- IKEv1 represents the IKE protocol version used when connecting to a remote peer.
- Step 6** Click **Create** to create this initiator version or click **Close** to discard any unsaved changes.
-

Clearing IKE Tunnels or Domains

To clear all the IKE tunnels or domains, follow these steps:

Procedure

-
- Step 1** Expand **Switches > Security**, and then select **IKE** in the Physical Attributes pane.
- Step 2** Click the **Tunnels** tab in the Information pane.
- You see the IKE tunnels.
- Step 3** Click the **Action** column and select **Clear** to clear the tunnel.
-

Refreshing SAs

Use the **crypto ike domain ipsec rekey IPv4-ACL-index** command to refresh the SAs after performing IKEv2 configuration changes.

To refresh the SAs after changing the IKEv2 configuration, follow these steps:

Procedure

-
- Step 1** Expand **Switches > Security**, and then select **IKE** in the Physical Attributes pane.
- Step 2** Click the **Pre-Shared AuthKey** tab in the Information pane.
- Step 3** Click **Refresh Values**.
-

Configuring Crypto

This sections includes the following topics:

Configuring Transform Sets


To configure transform sets, follow these steps:

Procedure

- Step 1** Expand **Switches > Security**, and then select **IPSec** in the Physical Attributes pane.
- Step 2** Click the **Transform Set** tab in the Information pane.
- Step 3** Click **Create Row**.

You see the Create IPSEC dialog box shown in [Figure 100: Create IPSEC, on page 731](#).

Figure 100: Create IPSEC



- Step 4** Select the switches that you want to create a transform set for in the Create Transform Set dialog box.
- Step 5** Assign a name and protocol for the transform set.
- Step 6** Select the encryption and authentication algorithm. See the IPsec Transform Configuration Parameters table to verify the allowed transform combinations.
- Step 7** Click **Create** to create the transform set or you click **Close**.

Creating Crypto Map Entries

To create mandatory crypto map entries, follow these steps:

Procedure

- Step 1** Expand **Switches > Security**, and then select **IPSEC** in the Physical Attributes pane.
- Step 2** Click the **CryptoMap Set Entry** tab.
- Step 3** (Optional) Click **Create Row** to create a crypto map entry.

- Step 4** Select the switch that you want to configure or modify. If you are creating a crypto map, set the setName and priority for this crypto map.
 - Step 5** Select the IPv4-ACL Profile and TransformSetIdList from the drop-down list for this crypto map.
 - Step 6** (Optional) Check the AutoPeer check box or set the peer address if you are creating a crypto map. See the [About the AutoPeer Option, on page 721](#).
 - Step 7** Choose the appropriate PFS selection. See the [About Perfect Forward Secrecy, on page 722](#).
 - Step 8** Supply the Lifetime and LifeSize. See the [About SA Lifetime Negotiation, on page 721](#).
 - Step 9** Click Create if you are creating a crypto map, or click Apply Changes if you are modifying an existing crypto map.
-

Setting the SA Lifetime

To set the SA lifetime for a specified crypto map entry, follow these steps:

Procedure

- Step 1** Expand Switches > Security and then select IPSEC in the Physical Attributes pane.
 - Step 2** Click the CryptoMap Set Entry tab.
 - Step 3** Scroll to the right half of the dialog box.
 - Step 4** Double-click and modify the value in the **Life Time(sec)** column.
 - Step 5** Click Apply Changes to save your changes.
-

Configuring Perfect Forward Secrecy

To configure the PFS value, follow these steps:

Procedure

- Step 1** Expand Switches > Security and then select IPSEC in the Physical Attributes pane.
 - Step 2** Click the CryptoMap Set Entry tab.
 - Step 3** From the drop-down list in the **PFS** column select the appropriate value.
 - Step 4** Click Apply Changes to save your changes.
-

Applying a Crypto Map Set

To apply a crypto map set to an interface, follow these steps:

Procedure

- Step 1** Expand Switches > Security, and then select IPSEC in the Physical Attributes pane.

- Step 2** Click the Interfaces tab.
- Step 3** Select the switch and interface you want to configure.
- Step 4** Enter the name of the crypto map that you want to apply to this interface in the CryptomapSetName field.
- Step 5** Click Create to apply the crypto map to the selected interface or click Close to exit the dialog box without applying the crypto map.

Configuring Global Lifetime Values

To configure global SA lifetimes, follow these steps:

Procedure

- Step 1** Choose Switches > Security, and then select IPSEC in the Physical Attributes pane.
- Step 2** You see the IPsec configuration in the Information pane.
- Step 3** Click the Global tab.
- Step 4** Double-click and edit the value in the **Life Time(sec)** column.
- Step 5** Click Apply Changes to save your changes.

Field Descriptions for IPsec

The following are the field descriptions for IPsec.

IPsec

Field	Description
Interface, CryptomapName	The binding of cryptomap sets to the interfaces of the managed entity.

IKE Global

Field	Description
RemIdentity	Displays the keepalive interval in seconds used by the IKE entity on the managed device with all the peers for the DOI corresponding to this conceptual row.
Key	Displays the type of keepalives to be used by the IKE entity on the managed device with all the peers for the DOI corresponding to this conceptual row.

IKE Pre-Shared AuthKey

Field	Description
KeepAliveInterval (sec)	The Phase 1 ID identity of the peer for which this pre-shared key is configured on the local entity.

Field	Description
IdentityType	The pre-shared authorization key used in authenticating the peer corresponding to this conceptual row.

IKE Policies

Field	Description
Priority	The priority of this ISAKMP policy entry. The policy with lower value would take precedence over the policy with higher value in the same DOI.
Encr	The encryption transform specified by this ISAKMP policy specification. The Internet Key Exchange (IKE) tunnels setup using this policy item would use the specified encryption transform to protect the ISAKMP PDUs.
Hash	The hash transform specified by this ISAKMP policy specification. The IKE tunnels setup using this policy item would use the specified hash transform to protect the ISAKMP PDUs.
Auth	The peer authentication method specified by this ISAKMP policy specification. If this policy entity is selected for negotiation with a peer, the local entity would authenticate the peer using the method specified by this object.
DHGroup	Specifies the Oakley group used for Diffie Hellman exchange in the Main Mode. If this policy item is selected to negotiate Main Mode with an IKE peer, the local entity chooses the group specified by this object to perform Diffie Hellman exchange with the peer.
Lifetime (sec)	Specifies the lifetime in seconds of the IKE tunnels generated using this policy specification.

IKE Initiator Version

Field	Description
Address	The address of the remote peer corresponding to this conceptual row. This object cannot be modified while the corresponding value of <code>cicIkeCfgInitiatorStatus</code> is equal to active.
Version	The IKE protocol version used when connecting to a remote peer specified in <code>cicIkeCfgInitiatorPAddr</code> . This object cannot be modified while the corresponding value of <code>cicIkeCfgInitiatorStatus</code> is equal to active.

IKE Tunnels

Field	Description
LocalAddress	The address of the local endpoint for the Phase-1 tunnel.
RemoteAddress	The address of the remote endpoint of the Phase-1 tunnel.
AuthMethod	The authentication method used in Phase-1 negotiations on the control tunnel corresponding to this conceptual row.

Field	Description
Action	The action to be taken on this tunnel. If clear, then this tunnel is cleared. If re-key, then re-keying is forced on this tunnel. The value none would be returned on doing read of this object.

IPSEC Global

Field	Description
Lifetime (sec)	The default lifetime (in seconds) assigned to an IPsec tunnel as a global policy (maybe overridden in specific cryptomap definitions).
Lifesize (KB)	The default life size in KB assigned to an IPsec tunnel as a global policy (unless overridden in cryptomap definition).

IPSEC Transform Set

Field	Description
Id	This is the sequence number of the transform set that uniquely identifies the transform set. Distinct transform sets must have distinct sequence numbers.
Protocol	Represents the suite of Phase-2 security protocols of this transform set.
ESP Encryption	Represents the transform used for ESP encryption.
ESP Authentication	Represents the transform used to implement integrity check with ESP protocol.
Mode	Represents the encapsulation mode of the transform set.

IPSEC CryptoMap Set Entry

Field	Description
IpFilter	Specifies an IP protocol filter to be secured using this cryptomap entry. When it has a value of zero-length string, it is not valid/applicable.
TransformSetIdList	The list of cipsXformSetId that are members of this CipsStaticCryptomapEntry. The value of this object is a concatenation of zero or more 4-octet strings, where each 4-octet string contains a 32-bit cipsXformSetId value in network byte order. A zero length string value means this list has no members.
AutoPeer	If true the destination address is taken as the peer address, while creating the tunnel.
Peer Address	The IP address of the peer to which this cryptomap entry is currently connected.
PFS	Identifies whether the tunnels instantiated due to this policy item should use Perfect Forward Secrecy (PFS) and if so, what group of Oakley they should use.
LifeTime	Specifies the lifetime of the IPsec Security Associations (SA) created using this IPsec policy entry.

Field	Description
Lifesize Value	Identifies the life size (maximum traffic in bytes that may be carried) of the IPsec SAs created using this IPsec policy entry. When a Security Association (SA) is created using this IPsec policy entry, its life size takes the value of this object.

IPSEC Interfaces

Field	Description
CryptomapName	The index of the static cryptomap table. The value of the string is the name string assigned by the NMS when defining a cryptomap set.
InterfaceList	Interfaces belong to the cryptomap.

IPSEC Tunnels

Field	Description
Local Address	The IP address of the local endpoint for the IPsec Phase-2 tunnel.
RemoteAddress	The type of the IP address of the remote endpoint for the IPsec Phase-2 tunnel.
ESP Encryption	The encryption algorithm used by the outbound security association of the IPsec Phase-2 tunnel.
ESP Encryption KeySize	The key size in bits of the negotiated key to be used with the algorithm denoted by ceipSecTunOutSaEncryptAlgo. For DES and 3DES the key size is respectively 56 and 168. For AES, this will denote the negotiated key size.
ESP Authentication	The authentication algorithm used by the inbound encapsulation security protocol (ESP) security association of the IPsec Phase-2 tunnel.
LifeSize (KB)	The negotiated life size of the IPSEC Phase-2 tunnel in kilobytes.
LifeTime (sec)	The negotiated lifetime of the IPSEC Phase-2 tunnel in seconds. If the tunnel was setup manually, the value of this MIB element should be 0.
Action	The status of the MIB table row.



CHAPTER 36

Configuring Port Security

- [Configuring Port Security, on page 737](#)

Configuring Port Security

All switches in the Cisco MDS 9000 Family provide port security features that reject intrusion attempts and report these intrusions to the administrator.

This chapter includes the following topics:

Information About Port Security

All switches in the Cisco MDS 9000 Family provide port security features that reject intrusion attempts and report these intrusions to the administrator.

Typically, any Fibre Channel device in a SAN can attach to any SAN switch port and access SAN services based on zone membership. Port security features prevent unauthorized access to a switch port in the Cisco MDS 9000 Family in the following ways:

- Login requests from unauthorized Fibre Channel devices (Nx ports) and switches (xE ports) are rejected.
- All intrusion attempts are reported to the SAN administrator through system messages.
- Configuration distribution uses the CFS infrastructure, and is limited to those switches that are CFS capable. Distribution is disabled by default.
- Configuring the port security policy requires the ENTERPRISE_PKG license (see the Cisco MDS 9000 Family NX-OS Licensing Guide).

This section includes the following topics:

Port Security Enforcement

To enforce port security, configure the devices and switch port interfaces through which each device or switch is connected, and activate the configuration.

- Use the port world wide name (pWWN) or the node world wide name (nWWN) to specify the Nx port connection for each device.
- Use the switch world wide name (sWWN) to specify the xE port connection for each switch.

Each Nx and xE port can be configured to restrict a single port or a range of ports.

Enforcement of port security policies are done on every activation and when the port tries to come up.

The port security feature uses two databases to accept and implement configuration changes.

- Configuration database—All configuration changes are stored in the configuration database.
- Active database—The database currently enforced by the fabric. The port security feature requires all devices connecting to a switch to be part of the port security active database. The software uses this active database to enforce authorization.

About Auto-Learning

You can instruct the switch to automatically learn (auto-learn) the port security configurations over a specified period. This feature allows any switch in the Cisco MDS 9000 Family to automatically learn about devices and switches that connect to it. Use this feature when you activate the port security feature for the first time as it saves tedious manual configuration for each port. You must configure auto-learning on a per-VSAN basis. If enabled, devices and switches that are allowed to connect to the switch are automatically learned, even if you have not configured any port access.

When auto-learning is enabled, learning happens only for the devices or interfaces that were not already logged into the switch. Learned entries on a port are cleaned up after you shut down that port if auto-learning is still enabled.

Learning does not override the existing configured port security policies. So, for example, if an interface is configured to allow a specific pWWN, then auto-learning will not add a new entry to allow any other pWWN on that interface. All other pWWNs will be blocked even in auto-learning mode.

No entries are learned for a port in the shutdown state.

When you activate the port security feature, auto-learning is also automatically enabled.



Note If you enable auto-learning before activating port security, you cannot activate until auto-learning is disabled.

Port Security Activation

By default, the port security feature is not activated in any switch in the Cisco MDS 9000 Family.

By activating the port security feature, the following apply:

- Auto-learning is also automatically enabled, which means:
 - From this point, auto-learning happens only for the devices or interfaces that were not logged into the switch.
 - You cannot activate the database until you disable auto-learning.
- All the devices that are already logged in are learned and are added to the active database.
- All entries in the configured database are copied to the active database.

After the database is activated, subsequent device login is subject to the activated port bound WWN pairs, excluding the auto-learned entries. You must disable auto-learning before the auto-learned entries become activated.

When you activate the port security feature, auto-learning is also automatically enabled. You can choose to activate the port security feature and disable auto-learning.



Tip If a port is shut down because of a denied login attempt, and you subsequently configure the database to allow that login, the port does not come up automatically. You must explicitly issue a no shutdown CLI command to bring that port back online.

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 PortChannel member.
- The configured database is empty but the active database is not.

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed by forcing the port security activation.

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.

Auto-learning Device Authorization

[Table 90: Authorized Auto-learning Device Requests](#), on page 739 summarizes the authorized connection conditions for device requests.

Table 90: 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

Condition	Device (pWWN, nWWN, sWWN)	Requests Connection to	Authorization
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 Scenarios

Assume that the port security feature is activated and the following conditions are specified in the active database:

- A pWWN (P1) is allowed access through interface fc1/1 (F1).
- A pWWN (P2) is allowed access through interface fc1/1 (F1).
- A nWWN (N1) is allowed access through interface fc1/2 (F2).
- Any WWN is allowed access through interface fc1/3 (F3).
- A nWWN (N3) is allowed access through any interface.
- A pWWN (P3) is allowed access through interface fc1/4 (F4).
- A sWWN (S1) is allowed access through interface fc1/10-13 (F10 to F13).
- A pWWN (P10) is allowed access through interface fc1/11 (F11).

Table 91: Authorization Results for Scenario , on page 740 summarizes the port security authorization results for this active database. The conditions listed refer to the conditions from [#unique_1636 unique_1636_Connect_42_tab_1000664](#).

Table 91: 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.
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.

Device Connection Request	Authorization	Condition	Reason
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.

About WWN Identification

If you decide to manually configure port security, be sure to adhere to the following guidelines:

- Identify switch ports by the interface or by the fWWN.
- Identify devices by the pWWN or by the nWWN.
- If an Nx port is allowed to log in to SAN switch port Fx, then that Nx port can only log in through the specified Fx port.
- If an Nx port's nWWN is bound to an Fx port WWN, then all pWWNs in the Nx port are implicitly paired with the Fx port.
- TE port checking is done on each VSAN in the allowed VSAN list of the trunk port.
- All PortChannel xE ports must be configured with the same set of WWNs in the same PortChannel.
- E port security is implemented in the port VSAN of the E port. In this case the sWWN is used to secure authorization checks.
- Once activated, the config database can be modified without any effect on the active database.
- By saving the running configuration, you save the configuration database and activated entries in the active database. Learned entries in the active database are not saved.

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, then the activation and auto-learning changes are consolidated and the behavior may change (see [Table 92: Scenarios for Activation and Auto-learning Configurations in Distributed Mode](#), on page 742).

Table 92: 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.	You activate the port security database and enable auto-learning.	configuration database = {A,B} active database = {A,B, C ³⁵ , D*}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled}
	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}
	You issue a commit.	Not applicable	configuration database = {A,B, E} active database = {A,B, E, C*, D*} pending database = empty
A and B exist in the configuration database, activation is not done and devices C,D are logged in.	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}
	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}
	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

³⁵ The * (asterisk):autolearned entries * (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 Interaction

[Table 93: Active and Configuration Port Security Databases](#), on page 743 lists the differences and interaction between the active and configuration databases.

Table 93: Active and Configuration Port Security Databases

Active Database	Configuration Database
Read-only.	Read-write.
Saving the configuration only saves the activated entries. Learned entries are not saved.	Saving the configuration saves all the entries in the configuration database.
Once activated, all devices that have already logged into the VSAN are also learned and added to the active database.	Once activated, the configuration database can be modified without any effect on the active database.
You can overwrite the active database with the configured database by activating the port security database. Forcing an activation may violate the entries already configured in the active database.	You can overwrite the configuration database with the active database.

Guidelines and Limitations

- Port security is only supported for Fibre Channel ports.

Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the active database.

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 2 K.



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.

Default Settings

[Table 94: Default Security Settings](#), on page 744 lists the default settings for all port security features in any switch.

Table 94: Default Security Settings

Parameters	Default
Auto-learn	Enabled if port security is enabled.
Port security	Disabled
Distribution	Disabled. Note Enabling distribution enables it on all VSANs in the switch.

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

To configure port security, using auto-learning and CFS distribution, follow these steps:

Procedure

-
- Step 1** Enable port security. See the [Enabling Port Security, on page 747](#).
 - Step 2** Enable CFS distribution. See the [Enabling Distribution, on page 752](#).
 - Step 3** Activate port security on each VSAN. This turns on auto-learning by default. See the [Activating Port Security, on page 747](#).
 - Step 4** Issue a CFS commit to copy this configuration to all switches in the fabric. See the [Committing the Changes, on page 753](#). At this point, all switches are activated, and auto-learning.
 - Step 5** Wait until all switches and all hosts are automatically learned.
 - Step 6** Disable auto-learn on each VSAN. See the [Disabling Auto-learning, on page 750](#).
 - Step 7** Issue a CFS commit to copy this configuration to all switches in the fabric. See the [Committing the Changes, on page 753](#). At this point, the auto-learned entries from every switch are combined into a static active database that is distributed to all switches.
 - Step 8** Copy the active database to the configure database on each VSAN. See the [Copying the Port Security Database, on page 753](#).
 - Step 9** Issue a CFS commit to copy this configuration to all switches in the fabric. See the [Committing the Changes, on page 753](#). This ensures that the configure database is the same on all switches in the fabric.
 - Step 10** Copy the running configuration to the startup configuration, using the fabric option. This saves the port security configure database to the startup configuration on all switches in the fabric.
-

Configuring Port Security with Auto-Learning without CFS

To configure port security using auto-learning without CFS, follow these steps:

Procedure

-
- | | |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Enable port security. See the Enabling Port Security, on page 747 . |
| Step 2 | Activate port security on each VSAN. This turns on auto-learning by default. |
| 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, on page 750 . |
| Step 5 | Copy the active database to the configure database on each VSAN. See the Copying the Port Security Database, on page 753 . |
| Step 6 | Copy the running configuration to the startup configuration. This saves the port security configure database to the startup configuration. |
| Step 7 | Repeat <code>#unique_1651 unique_1651_Connect_42__1001237</code> through <code>#unique_1651 unique_1651_Connect_42__1001252</code> for all switches in the fabric. |
-

Configuring Port Security with Manual Database Configuration

To configure port security and manually configure the port security database, follow these steps:

Procedure

-
- | | |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Enable port security. See the Enabling Port Security, on page 747 . |
| Step 2 | Manually configure all port security entries into the configure database on each VSAN. See the Configuring Port Security with Manual Database Configuration, on page 745 . |
| Step 3 | Activate port security on each VSAN. This turns on auto-learning by default. See the Activating Port Security, on page 747 . |
| Step 4 | Disable auto-learn on each VSAN. See the Disabling Auto-learning, on page 750 . |
| Step 5 | Copy the running configuration to the startup configuration. This saves the port security configure database to the startup configuration. |
| Step 6 | Repeat <code>#unique_1652 unique_1652_Connect_42__1001270</code> through <code>#unique_1652 unique_1652_Connect_42__1001284</code> for all switches in the fabric. |
-

Configuring Port Security Using the Configuration Wizard

The Port Security Configuration wizard provides step-by-step procedures for setting up the Port Security Policy for a selected VSAN. The Port Security Configuration wizard also supports the central management through CFS, making it possible to complete the entire configuration at one place.

The wizard automatically conducts few essential operations. For example, if you want central management, the wizard conducts operations to check CFS capability, enable CFS, and issue CFS commit at the proper stages.

To manage security at a particular port, you do not need to run through the wizard to configure the port security policy from the VSAN wide, but you can directly edit accesses on the port itself. This operation can be done through the Port Binding dialog box. If the port's belonging switch has not enabled port security yet, the dialog box enables security first. If the port security is enabled, the dialog box will edit the policy database based on user operations.

CFS should be enabled on all switches in the VSAN. A CFS master switch is selected to do all configurations. All changes will be distributed to the VSAN through the CFS commit command.


To configure port security, follow these steps:

Before you begin

The following are the prerequisites:

- Enable port security on the switch.
- Define port security policy either manually by editing bound devices or switches or ports or by using autolearning.
- Activate port security policy.
- Ensure that activated and configured databases are synchronized through copy.
- Copy the activated database to be the startup configuration.

Procedure

	Command or Action	Purpose
Step 1	Before launching the Port Security Setup Wizard, DCNM-SAN checks the CFS capability of the switches in the VSAN.	<p>If VSAN context is not available, the wizard prompts to select VSAN as shown in Figure 101: Select VSAN Window, on page 746</p> <p>Figure 101: Select VSAN Window</p> 
Step 2	Select the VSAN from the list and click OK . Do the following in the Select Master Switch page:	<ul style="list-style-type: none"> • Select the required master switch. • Select Automatically learn all logged in ports in VSAN to Autolearn port configuration.
Step 3	Click Next to proceed.	<p>You see the Edit and Activate Configuration page.</p> <p>Note From Cisco NX-OS Release 5.2, devices can bind to vFC interfaces.</p>
Step 4	Click Insert to create port binding.	<p>Note When interfaces are inserted for binding, vFC ports can be selected.</p> <p>Two types of port binding can be created using the Insert Port Security Devices dialog box:</p> <ul style="list-style-type: none"> • Port WWN-pWWN bound to an interface WWN.

	Command or Action	Purpose
		• Switch -Switch WWN bound to an interface. (Mainly useful for ISL binding).
Step 5	Select the type of port binding by clicking the radio buttons and enter the supporting values.	
Step 6	Click OK .	
Step 7	Click Close to exit the Insert Port Security window.	Note To delete an entry in the Edit and Activate Configuration page of the wizard, select the entry and click the Delete button.
Step 8	Click Finish to complete the Port Security Configuration for the selected switch.	

Enabling Port Security

By default, the port security feature is disabled in all switches in the Cisco MDS 9000 Family.

To enable port security, follow these steps:

Procedure

-
- Step 1** Expand a VSAN, and then select Port Security in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the CFS tab.
- 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.
- Step 6** Set the Command column to enable for each switch in the VSAN.
- Step 7** Click the CFS tab and set the Command column to commit on all participating switches in the VSAN.
- Step 8** Click Apply Changes to distribute the enabled port security to all switches in the VSAN.
-

Activating Port Security

To activate port security, follow these steps:

Procedure

-
- Step 1** Expand a VSAN and select Port Security in the Logical Domains pane.

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

- Step 2** Click the Actions tab.
- Step 3** Click in the Action column under Activation, next to the switch or VSAN on which you want to activate port security. You see a drop-down menu with options.
- 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 DCNM-SAN or Apply in Device Manager to save these changes.

Note If required, you can disable auto-learning (see the [Disabling Auto-learning, on page 750](#)).

Activating the Port Security Forcefully

If the port security activation request is rejected, you can force the activation.



Note An activation using the **force** option can log out existing devices if they violate the active database.

You can view missing or conflicting entries using the **port-security database diff active vsan** command in EXEC mode.

To forcefully activate the port security database, follow these steps:

Procedure

- Step 1** Expand a VSAN and select Port Security in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the Actions tab.
- Step 3** Click in the Action column under Activation, next to the switch or VSAN on which you want to activate port security and select the **forceactivate** option.
- Step 4** Set the Action field you want for that switch.
- Step 5** Click the CFS tab and set the command column to commit on all participating switches in the VSAN.
- Step 6** Click Apply Changes in DCNM-SAN or Apply in Device Manager to save these changes.

Reactivating the Database



Tip If auto-learning is enabled, and you cannot activate the database, you will not be allowed to proceed without the force option until you disable auto-learning..

To reactivate the port security database, follow these steps:

Procedure

- Step 1** Disable auto-learning.
- Step 2** Copy the active database to the configured database.
- Tip** If the active database is empty, you cannot perform this step.
- Step 3** Make the required changes to the configuration database.
- Step 4** Activate the database.
-

Copying an Active Database to the Config Database

To copy the active database to the config database, follow these steps:

Procedure

- Step 1** Expand a VSAN and select Port Security in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the Actions tab.
You see the switches for that VSAN.
- Step 3** Check the CopyActive ToConfig check box next to the switch for which you want to copy the database.
The active database is copied to the config database when the security setting is activated.
- Step 4** Uncheck the CopyActive ToConfig check box if you do not want the database copied when the security setting is activated.
- Step 5** Click the CFS tab and set the command column to commit on all participating switches in the VSAN.
- Step 6** Click Apply Changes to save these changes or click Undo Changes to discard any unsaved changes.
-

Configuring Auto-learning

This section contains the following topics:

Enabling Auto-learning

To enable auto-learning, follow these steps:

Procedure

- Step 1** Expand a VSAN and select Port Security in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.

- Step 2** Click the Actions tab.
- Step 3** Click in the Action column under Activation, next to the switch or VSAN on which you want to activate port security. You see a drop-down menu with the following options:
- **activate**—Valid port security settings are activated.
 - **activate (TurnLearningOff)**—Valid port security settings are activated and auto-learn turned off.
 - **forceActivate**—Activation is forced.
 - **forceActivate(TurnLearningOff)**—Activation is forced and auto-learn is turned off.
 - **deactivate**—All currently active port security settings are deactivated.
 - **NoSelection**—No action is taken.
- Step 4** Select one of the port security options for that switch.
- Step 5** Check the AutoLearn check box for each switch in the VSAN to enable auto-learning.
- Step 6** Click the Apply Changes icon to save these changes.
-

Disabling Auto-learning

To disable auto-learning, follow these steps:

Procedure

-
- Step 1** Expand a VSAN and select Port Security in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the Actions tab.
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.
-

Configuring Port Security Manually

This section includes the following topics:

Task Flow for Configuring Port Security

Follow these steps to configure port security on any switch in the Cisco MDS 9000 Family:

Procedure

-
- Step 1** Identify the WWN of the ports that need to be secured.

- Step 2** Secure the fWWN to an authorized nWWN or pWWN.
 - Step 3** Activate the port security database.
 - Step 4** Verify your configuration.
-

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, follow these steps:

Procedure

- Step 1** Expand a VSAN and select 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.
 - 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, follow these steps:

Procedure

- Step 1** Expand a VSAN and select 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 the row you want to delete.
- Step 4** Click Delete Row.
You see the confirmation dialog box.
- Step 5** Click Yes to delete the row, or click No to close the confirmation dialog box without deleting the row.

Step 6 Click the Apply Changes icon to save these changes.

Configuring Port Security 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 35, “Configuring IPsec Network Security.”](#)).

This section includes the following topics:

Enabling Distribution

For example, if you activate port security, follow up by disabling auto-learning, and commit the changes in the pending database, then the net result of your actions is the same as issuing a **port-security activate vsan vsan-id no-auto-learn** command.

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, on page 741](#).



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.

To enable distribution, follow these steps:

Procedure

- Step 1** Expand a VSAN and select Port Security in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the Control tab.
You see the switches for that VSAN.
- Step 3** In the Command column, select enable or disable from the drop-down menu.
- 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.

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.

Discarding the Changes

If you discard (abort) the changes made to the pending database, the configuration remains unaffected and the lock is released.

Interacting with the Database

**Note**

You can overwrite the configuration database with the active database using the **port-security database copy vsan** command. The **port-security database diff active vsan** command in EXEC mode lists the differences between the active database and the configuration database.

This section includes the following topics:

Copying the Port Security Database

Use the **port-security database copy vsan** command to copy from the active to the configured database. If the active database is empty, this command is not accepted.

```
switch# port-security database copy vsan 1
```

Use the **port-security database diff active vsan** command to view the differences between the active database and the configuration database. This command can be used when resolving conflicts.

```
switch# port-security database diff active vsan 1
```

Use the **port-security database diff config vsan** command to obtain information on the differences between the configuration database and the active database.

```
switch# port-security database diff config vsan 1
```

**Tip**

We recommend that you copy the active database to the configuration database issue the **port-security database copy vsan** command after disabling auto-learning. This action ensures that the configuration database is in sync 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.

copy the active database to the configuration database

To copy the active database to the configuration database, follow these steps:

Procedure

-
- Step 1** Expand a **Fabric**, expand a **VSAN**, and then select **Port Security** in the Logical Domains pane.
 - Step 2** Click the **Actions** tab. You see all the configuration databases.
 - Step 3** Select the appropriate configuration database and check the **Copy Active to Config** check box.
 - Step 4** Click the **Apply Changes** icon to save your changes.
-

View differences between Active DB and Configuration DB

To view the differences between the active database and the configuration database, follow these steps:

Procedure

-
- Step 1** Expand a **Fabric**, expand a **VSAN**, and then select **Port Security** in the Logical Domains pane.
You see the Port Security information in the Information pane.
 - Step 2** Click the **Database Differences** tab. You see all the configuration databases.
 - Step 3** Select the appropriate configuration database. Select the **Active** or **Config** option to compare the differences between the selected database and the active or configuration database.
 - Step 4** Click 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 **port-security commit** command is required to actually delete the database.

Use the **no port-security database vsan** command in configuration mode to delete the configured database for a specified VSAN

```
switch(config)# no port-security database vsan 1
```

To delete a port security database, follow these steps:

Procedure

-
- Step 1** Expand a **Fabric**, expand a **VSAN**, and then select **Port Security** in the Logical Domains pane.
You see the Port Security information in the Information pane.
 - Step 2** Click the **Config Database** tab. You see all the configuration databases.

- Step 3** Select the appropriate configuration database and click the **Delete Row** button.
- Step 4** Click **Yes** if you want to delete the configuration database.

Cleaning the Port Security Database

Use the **clear port-security statistics vsan** command to clear all existing statistics from the port security database for a specified VSAN.

```
switch# clear port-security statistics vsan 1
```

Use the **clear port-security database auto-learn interface** command to clear any learned entries in the active database for a specified interface within a VSAN.

```
switch# clear port-security database auto-learn interface fc1/1 vsan 1
```

Use the **clear port-security database auto-learn vsan** command to clear any learned entries in the active database for the entire VSAN.

```
switch# clear port-security database auto-learn vsan 1
```



Note The **clear port-security database auto-learn** and **clear port-security statistics** commands are only relevant to the local switch and do not acquire locks. Also, learned entries are only local to the switch and do not participate in distribution.

Use the **port-security clear vsan** command to clear the pending session in the VSAN from any switch in the VSAN.

```
switch# clear
port-security session vsan 5
```

To clear all existing statistics from the port security database for a specified VSAN, follow these steps:

Procedure

- Step 1** Expand a **Fabric**, expand a **VSAN**, and then select **Port Security** in the Logical Domains pane. You see the Port Security information in the Information pane.
- Step 2** Click the **Statistics** tab. You see all the configuration databases.
- Step 3** Select the appropriate configuration database and check the **Clear** option.
- Step 4** Click the **Apply Changes** icon to save your changes.

Cleaning the Port Security Database

To clear any learned entries in the active database for a specified interface within a VSAN, follow these steps:

Procedure

- Step 1** Expand a **Fabric**, expand a **VSAN**, and then select **Port Security** in the Logical Domains pane.
You see the Port Security information in the Information pane.
- Step 2** Select the **Actions** tab. You see all the configuration databases.
- Step 3** Select the appropriate configuration database and check the **AutoLearn** option.
- Step 4** Click the **Apply Changes** icon to save your changes.
-

What to do next



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.

Displaying Activated Port Security Settings

To display active port security settings, follow these steps:

Procedure

- Step 1** Expand a VSAN and select Port Security in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the Active Database tab.
You see the active port security settings for that VSAN.
-

Displaying Port Security Statistics

To display port security statistics, follow these steps:

Procedure

- Step 1** Expand a VSAN and select Port Security in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the Statistics tab.
You see the port security statistics for that VSAN.
-

Displaying Port Security Violations

Port violations are invalid login attempts (for example, login requests from unauthorized Fibre Channel devices). You can display a list of these attempts on a per-VSAN basis.

To display port security violations, follow these steps:

Procedure

-
- Step 1** Expand a VSAN and select Port Security in the Logical Domains pane.
You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the Violations tab. You see the port security violations for that VSAN.
-

Field Descriptions for Port Security

The following are the field descriptions for port security.

Port Security Actions

Field	Description
Activation	
Action	<ul style="list-style-type: none"> • activate—Results in the valid port bindings on this VSAN/VLAN being activated. • activate (Turn LearningOff)—Results in the valid port bindings on this VSAN/VLAN being activated and copied to the active database and will also result in auto learn being turned off on this VSAN/VLAN, once the activation is complete. • force activate—Results in forced activation, even if there are errors during activation and the activated port bindings will be copied to the active database. • force activate (Turn Learning Off)—Results in forced activation along with turning auto learn off after activation and the activated port bindings will be copied to the active database. • deactivate—Results in deactivation of currently activated valid port bindings (if any), on this VSAN/VLAN. Currently active entries (if any), which would have been present in the active database, will be removed. • Activation will not be allowed on a VSAN if auto-learn is enabled on that VSAN
Enabled	The state of activation on this VSAN/VLAN. If true, then an activation has been attempted as the most recent operation on this VSAN/VLAN. If false, then an activation has not been attempted as the most recent operation on this VSAN/VLAN.
Result	Indicates the outcome of the most recent activation/deactivation.
Last Change	When the valid port bindings on this VSAN/VLAN were last activated. If the last activation took place prior to the last re-initialization of the agent, then this value will be N/A.

Field	Description
CopyActiveToConfig	If enabled, results in the active port binding database to be copied on to the configuration database on this VSAN/VLAN. Note that the learned entries are also copied.
AutoLearn	Helps to learn the valid port binding configuration of devices/ports logged into the local device on all its ports and populate the above active database with the same. This mechanism of learning the configuration of devices/ports logged into the local device over a period of time and populating the configuration is a convenience mechanism for users. If enabled on a particular VSAN, all subsequent logins (FLOGIs) on that VSAN will be populated in the enforced port binding database, provided it is not in conflict with existing enforced port bindings on that VSAN. When disabled, the mechanism of learning is stopped. The learned entries will however be in the active database.
Clear AutoLearned	<ul style="list-style-type: none"> • Clear VSAN results in port bind auto-learned entries being cleared on this VSAN. • Clear Interface(s) results in port bind auto-learned entries being cleared on the interface specified on this VSAN.
Action	
Interface	Specifies the interface(s) on which the port bind auto-learned entries need to be cleared.

Port Security Config Database

Field	Description
Interface or fWWN	Represents the address of the port on the local device through which the device specified can FLOGI. <ul style="list-style-type: none"> • If fwwn, then the value is the fabric WWN of a port on the local device. • If intfIndex, then a port on the local device is being represented by its interface. • If wildCard, then it represents a wild-card entry. The wild-card represents any port on the local device.
Type	The mechanism to identify a switch port.
WWN	Represents the logging-in device address.
Available Interface	Displays the available interface. The interfaces available are: <ul style="list-style-type: none"> • Fibre Channel • PortChannel • Ethernet PortChannel • VFC

Port Security Active Database

Field	Description
Interface or fWWN	The address of a port on the local device.

Field	Description
Type	The mechanism to identify a switch port. <ul style="list-style-type: none"> • fwwn— The local switch port is identified by Fabric WWN(fWWN). • intfIndex— The local switch port is identified by ifIndex. • wildCard— Wild card (any switch port on local device).
WWN	Represents the logging-in device address.
IsLearnt	Indicates if this entry is a learned entry or not.

Port Security Database Differences

Field	Description
CompareWith	Specifies the database for the comparison. <ul style="list-style-type: none"> • configDb— Compares the configuration database with respect to active database on this VSAN/VLAN. So, the active database will be the reference database and the results of the difference operation will be with respect to the active database. • activeDb— Compares the active database with respect to configuration database on this VSAN/VLAN. So, the configuration database will be the reference database and the results of the difference operation will be with respect to the configuration database.
VSANId	The ID of the VSAN to compare against.
Interface/fWWN	The address of a port on the local device.
Type	The mechanism to identify a switch port. <ul style="list-style-type: none"> • fwwn— The local switch port is identified by Fabric WWN(fWWN). • intfIndex— The local switch port is identified by ifIndex. • wildCard— Wild card (any switch port on local device).
WWN	Represents the logging in device address.
Reason	Indicates the reason for the difference between the databases being compared, for this entry.

Port Security Violations

Field	Description
Interface	The fWWN of the port on the local device where the login was denied.
End Device	The pWWN of the device that was denied FLOGI on one of the local device's ports.
Or Switch	The sWWN of the device (if the device happens to be a switch), that was denied entry on one of the local device's ports.

Field	Description
Time	When the login denial took place.
Count	The number of times this particular pWWN/nWWN or sWWN has been denied login on this particular local interface.

Port Security Statistics

Field	Description
AllowedLogins	The number of FLOGI requests that have been allowed on this VSAN/VLAN.
DeniedLogins	The number of FLOGI requests that have been denied on this VSAN/VLAN.
Clear	When set to clear, it results in port bind statistic counters being cleared on this VSAN/VLAN.



CHAPTER 37

Configuring Fabric Binding

- [Configuring Fabric Binding, on page 761](#)

Configuring Fabric Binding

This chapter describes the fabric binding feature provided in the Cisco MDS 9000 Family of directors and switches. It includes the following sections:

This chapter includes the following topics:

Information About Fabric Binding

The fabric binding feature ensures ISLs are only enabled between specified switches in the fabric binding configuration. Fabric binding is configured on a per-VSAN basis.

This feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. It uses the Exchange Fabric Membership Data (EFMD) protocol to ensure that the list of authorized switches is identical in all switches in the fabric.

This section includes the following topics:

Port Security Versus Fabric Binding

Port security and fabric binding are two independent features that can be configured to complement each other. [Table 95: Fabric Binding and Port Security Comparison](#), on page 761 compares the two features.

Table 95: Fabric Binding and Port Security Comparison

Fabric Binding	Port Security
Uses a set of sWWNs and a persistent domain ID.	Uses pWWNs and nWWNs or fWWNs and sWWNs.
Binds the fabric at the switch level.	Binds devices at the interface level.

Fabric Binding	Port Security
Authorizes only the configured sWWN stored in the fabric binding database to participate in the fabric.	Allows a preconfigured set of Fibre Channel devices to logically connect to a SAN ports. The switch port, identified by a WWN or interface number, connects to a Fibre Channel device (a host or another switch), also identified by a WWN. By binding these two devices, you lock these two ports into a group (or list).
Requires activation on a per VSAN basis.	Requires activation on a per VSAN basis.
Allows specific user-defined switches that are allowed to connect to the fabric, regardless of the physical port to which the peer switch is connected.	Allows specific user-defined physical ports to which another device can connect.
Does not learn about switches that are logging in.	Learns about switches or devices that are logging in if learning mode is enabled.
Cannot be distributed by CFS and must be configured manually on each switch in the fabric.	Can be distributed by CFS.

Port-level checking for 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. In a FICON VSAN, the fabric binding feature requires all sWWNs connected to a switch and their persistent domain IDs to be part of the fabric binding active database. In a Fibre Channel VSAN, only the sWWN is required; the domain ID is optional.



Note

All switches in a Fibre Channel VSAN using fabric binding must be running Cisco MDS SAN-OS Release 3.0(1) and NX-OS Release 4.1(1b) or later.

Licensing Requirements for Fabric Binding

Fabric binding requires that you install either the **MAINFRAME_PKG** license or the **ENTERPRISE_PKG** license on your switch.

See the Cisco MDS 9000 Family NX-OS Licensing Guide for more information on license feature support and installation.

Default Settings

[Table 96: Default Fabric Binding Settings , on page 763](#) lists the default settings for the fabric binding feature.

Table 96: Default Fabric Binding Settings

Parameters	Default
Fabric binding	Disabled

Configuring Fabric Binding

To configure fabric binding in each switch in the fabric, follow these steps:

Procedure

-
- Step 1** Enable the fabric configuration feature.
 - Step 2** Configure a list of sWWNs and their corresponding domain IDs for devices that are allowed to access the fabric.
 - Step 3** Activate the fabric binding database.
 - Step 4** Copy the fabric binding active database to the fabric binding config database.
 - Step 5** Save the fabric binding configuration.
 - Step 6** Verify the fabric binding configuration.
-

Enabling Fabric Binding

The fabric binding feature must be enabled in each switch in the fabric that participates in the fabric binding. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family. The configuration and verification commands for the fabric binding feature are only available when fabric binding is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.



CHAPTER 38

Configuring FCIP

- [Configuring FCIP, on page 765](#)

Configuring FCIP

Cisco MDS 9000 Family IP storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch can connect separated SAN islands using Fibre Channel over IP (FCIP).




Note FCIP is supported on the MDS 9222i switch, MSM-18/4 module, MDS 9216i switch, MPS-14/2 module, 16-Port Storage Services Node (SSN-16), and IPS modules on MDS 9200 Series directors.

This chapter includes the following topics:

Information About FCIP

The Fibre Channel over IP Protocol (FCIP) is a tunneling protocol that connects geographically distributed Fibre Channel storage area networks (SAN islands) transparently over IP local area networks (LANs), metropolitan area networks (MANs), and wide area networks (WANs). The switch can connect separated SAN islands using Fibre Channel over IP (FCIP) (see [Figure 102: Fibre Channel SANs Connected by FCIP, on page 765](#)).

Figure 102: Fibre Channel SANs Connected by FCIP



FCIP uses TCP as a network layer transport. The DF bit is set in the TCP header.

**Note**

For more information about FCIP protocols, refer to the IETF standards for IP storage at <http://www.ietf.org>. Also refer to Fibre Channel standards for switch backbone connection at <http://www.t11.org> (see FC-BB-2).

This section includes the following topics:

FCIP Concepts

To configure IPS modules or MPS-14/2 modules for FCIP, you should have a basic understanding of the following concepts:

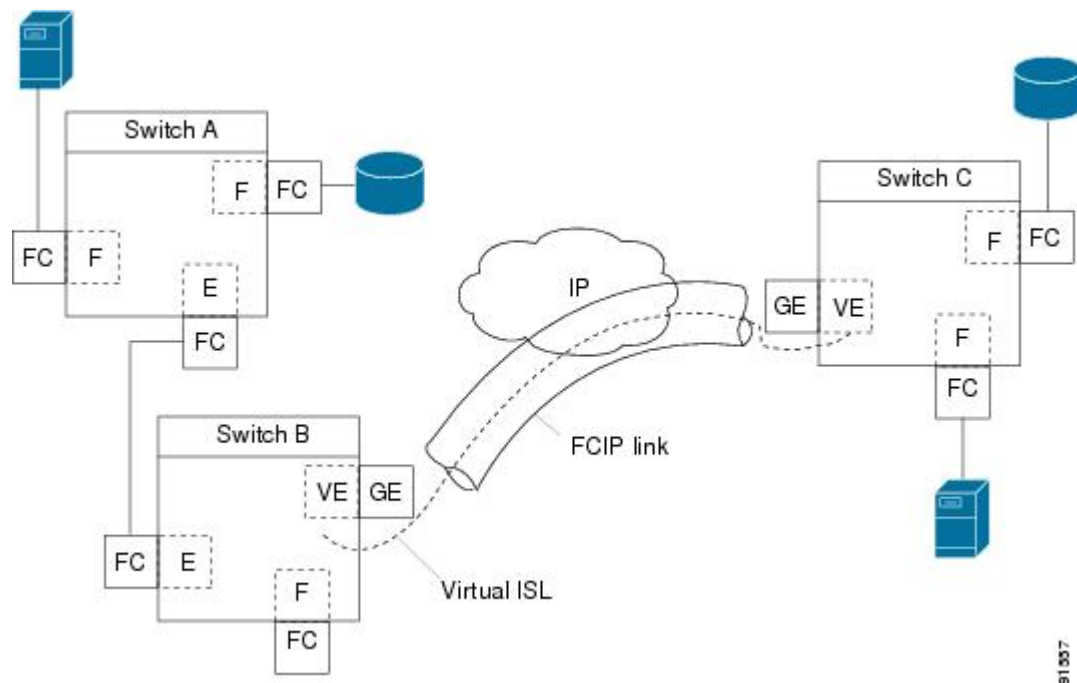
FCIP and VE Ports

[Figure 103: FCIP Links and Virtual ISLs, on page 766](#) describes the internal model of FCIP with respect to Fibre Channel Inter-Switch Links (ISLs) and Cisco's extended ISLs (EISLs).

FCIP virtual E (VE) ports behave exactly like standard Fibre Channel E ports, except that the transport in this case is FCIP instead of Fibre Channel. The only requirement is for the other end of the VE port to be another VE port.

A virtual ISL is established over an FCIP link and transports Fibre Channel traffic. Each associated virtual ISL looks like a Fibre Channel ISL with either an E port or a TE port at each end (see [Figure 103: FCIP Links and Virtual ISLs, on page 766](#)).

Figure 103: FCIP Links and Virtual ISLs



See the [Configuring B Ports, on page 791](#) for more information.

FCIP Links

FCIP links consist of one or more TCP connections between two FCIP link endpoints. Each link carries encapsulated Fibre Channel frames.

When the FCIP link comes up, the VE ports at both ends of the FCIP link create a virtual Fibre Channel (E)ISL and initiate the E port protocol to bring up the (E)ISL.

By default, the FCIP feature on any Cisco MDS 9000 Family switch creates two TCP connections for each FCIP link:

- One connection is used for data frames.
- The other connection is used only for Fibre Channel control frames, that is, switch-to-switch protocol frames (all Class F). This arrangement provides low latency for all control frames.

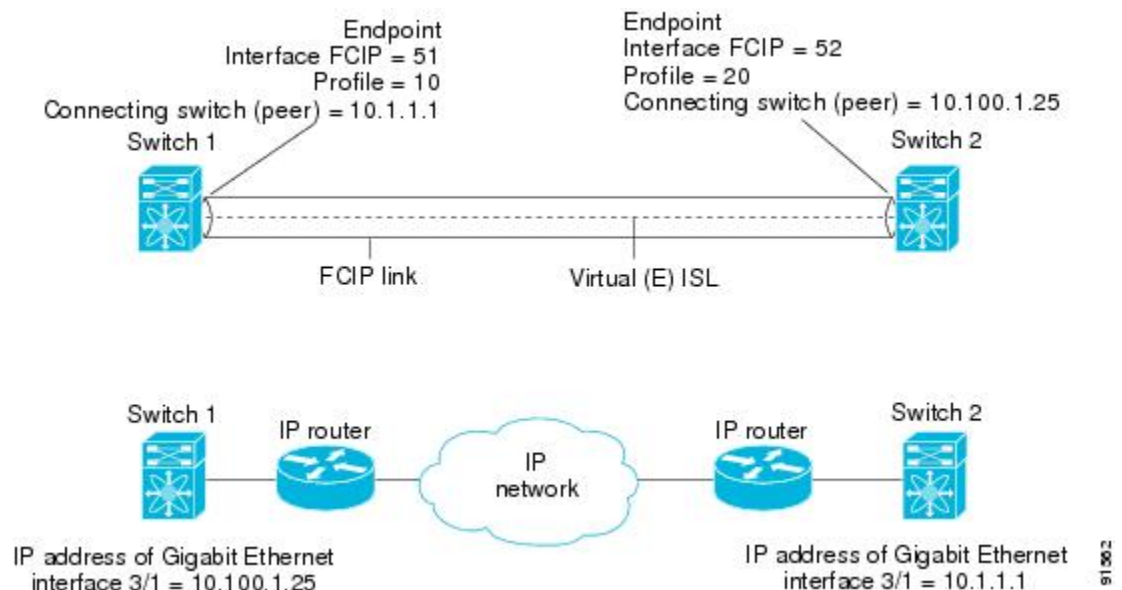
To enable FCIP on the IPS module or MPS-14/2 module, an FCIP profile and FCIP interface (interface FCIP) must be configured.

The FCIP link is established between two peers, the VE port initialization behavior is identical to a normal E port. This behavior is independent of the link being FCIP or pure Fibre Channel, and is based on the E port discovery process (ELP, ESC).

Once the FCIP link is established, the VE port behavior is identical to E port behavior for all inter-switch communication (including domain management, zones, and VSANs). At the Fibre Channel layer, all VE and E port operations are identical.

When two FCIP link endpoints are created, an FCIP link is established between the two IPS modules or MPS-14/2 modules. To create an FCIP link, assign a profile to the FCIP interface and configure the peer information. The peer IP switch information initiates (creates) an FCIP link to that peer switch (see [Figure 104: Assigning Profiles to Each Gigabit Ethernet Interface, on page 767](#)).

Figure 104: Assigning Profiles to Each Gigabit Ethernet Interface




FCIP Profiles

The FCIP profile contains information about the local IP address and TCP parameters. The profile defines the following information:

- The local connection points (IP address and TCP port number)
- The behavior of the underlying TCP connections for all FCIP links that use this profile

The FCIP profile's local IP address determines the Gigabit Ethernet port where the FCIP links terminate (see [Figure 105: FCIP Profile and FCIP Links, on page 768](#)).


Figure 105: FCIP Profile and FCIP Links



You must assign a local IP address of a Gigabit Ethernet interface or subinterface to the FCIP profile to create an FCIP profile. You can assign IPv4 or IPv6 addresses to the interfaces.

[Figure 106: Assigning Profiles to Each Gigabit Ethernet Interface, on page 768](#) shows an example configuration.

Figure 106: Assigning Profiles to Each Gigabit Ethernet Interface



FCIP Interfaces

The FCIP interface is the local endpoint of the FCIP link and a VE port interface. All the FCIP and E port parameters are configured in context to the FCIP interface.

The FCIP parameters consist of the following:

- The FCIP profile determines which Gigabit Ethernet port initiates the FCIP links and defines the TCP connection behavior.
- Peer information.
- Number of TCP connections for the FCIP link.
- E port parameters—trunking mode and trunk allowed VSAN list.


FCIP High-Availability Solutions

The following high-availability solutions are available for FCIP configurations:

Fibre Channel PortChannels

Figure 107: PortChannel-Based Load Balancing, on page 769 provides an example of a PortChannel-based load-balancing configuration. To perform this configuration, you need two IP addresses on each SAN island. This solution addresses link failures.

Figure 107: PortChannel-Based Load Balancing




The following characteristics set Fibre Channel PortChannel solutions apart from other solutions:

- The entire bundle is one logical (E)ISL link.
- All FCIP links in the PortChannel should be across the same two switches.
- The Fibre Channel traffic is load balanced across the FCIP links in the PortChannel.

FSPF

Figure 108: FSPF-Based Load Balancing, on page 769 displays a FSPF-based load balancing configuration example. This configuration requires two IP addresses on each SAN island, and addresses IP and FCIP link failures.

Figure 108: FSPF-Based Load Balancing



The following characteristics set FSPF solutions apart from other solutions:


- Each FCIP link is a separate (E)ISL.
- The FCIP links can connect to different switches across two SAN islands.

- The Fibre Channel traffic is load balanced across the FCIP link.

VRRP

Figure 109: VRRP-Based High Availability, on page 770 displays a Virtual Router Redundancy Protocol (VRRP)-based high availability FCIP configuration example. This configuration requires at least two physical Gigabit Ethernet ports connected to the Ethernet switch on the island where you need to implement high availability using VRRP.

Figure 109: VRRP-Based High Availability




The following characteristics set VRRP solutions apart from other solutions:

- If the active VRRP port fails, the standby VRRP port takes over the VRRP IP address.
- When the VRRP switchover happens, the FCIP link automatically disconnects and reconnects.
- This configuration has only one FCIP (E)ISL link.

Ethernet PortChannels

Figure 110: Ethernet PortChannel-Based High Availability, on page 770 displays an Ethernet PortChannel-based high-availability FCIP example. This solution addresses the problem caused by individual Gigabit Ethernet link failures.

Figure 110: Ethernet PortChannel-Based High Availability



The following characteristics set Ethernet PortChannel solutions apart from other solutions:


- The Gigabit Ethernet link-level redundancy ensures a transparent failover if one of the Gigabit Ethernet links fails.
- Two Gigabit Ethernet ports in one Ethernet PortChannel appear like one logical Gigabit Ethernet link.
- The FCIP link stays up during the failover.

Ethernet PortChannels and Fibre Channel PortChannels

Ethernet PortChannels offer link redundancy between the Cisco MDS 9000 Family switch's Gigabit Ethernet ports and the connecting Ethernet switch. Fibre Channel PortChannels also offer (E)ISL link redundancy between Fibre Channel switches. FCIP is an (E)ISL link and is only applicable for a Fibre Channel PortChannel. Beneath the FCIP level, an FCIP link can run on top of an Ethernet PortChannel or on one Gigabit Ethernet port. This link is totally transparent to the Fibre Channel layer.

An Ethernet PortChannel restriction only allows two contiguous IPS ports, such as ports 1–2 or 3–4, to be combined in one Ethernet PortChannel (see Chapter 42, “*Configuring IP Storage*” for more information). This restriction only applies to Ethernet PortChannels. The Fibre Channel PortChannel (to which FCIP link can be a part of) does not have a restriction on which (E)ISL links can be combined in a Fibre Channel PortChannel as long as it passes the compatibility check. The maximum number of Fibre Channel ports that can be put into a Fibre Channel PortChannel is 16 (see [Figure 111: PortChannels at the Fibre Channel and Ethernet Levels](#) on page 771).

Figure 111: PortChannels at the Fibre Channel and Ethernet Levels



To configure Fibre Channel PortChannels, see the Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide Interfaces Configuration Guide, Cisco DCNM for SAN.

To configure Ethernet PortChannels, see the High Availability and Redundancy Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS High Availability and Redundancy Configuration Guide.

FCIP Profile Configuration

A basic FCIP configuration uses the local IP address to configure the FCIP profile. In addition to the local IP address and the local port, you can specify other TCP parameters as part of the FCIP profile configuration.

Peers

All the FCIP and E port parameters are configured in context to the FCIP interface. To create an FCIP link, assign a profile to the FCIP interface and configure the peer information. The peer IP switch information

initiates (creates) an FCIP link to that peer switch. The basic FCIP configuration uses the peer's IP address to configure the peer information. You can establish an FCIP link with the peer using the Peer IP address option. This option configures both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.

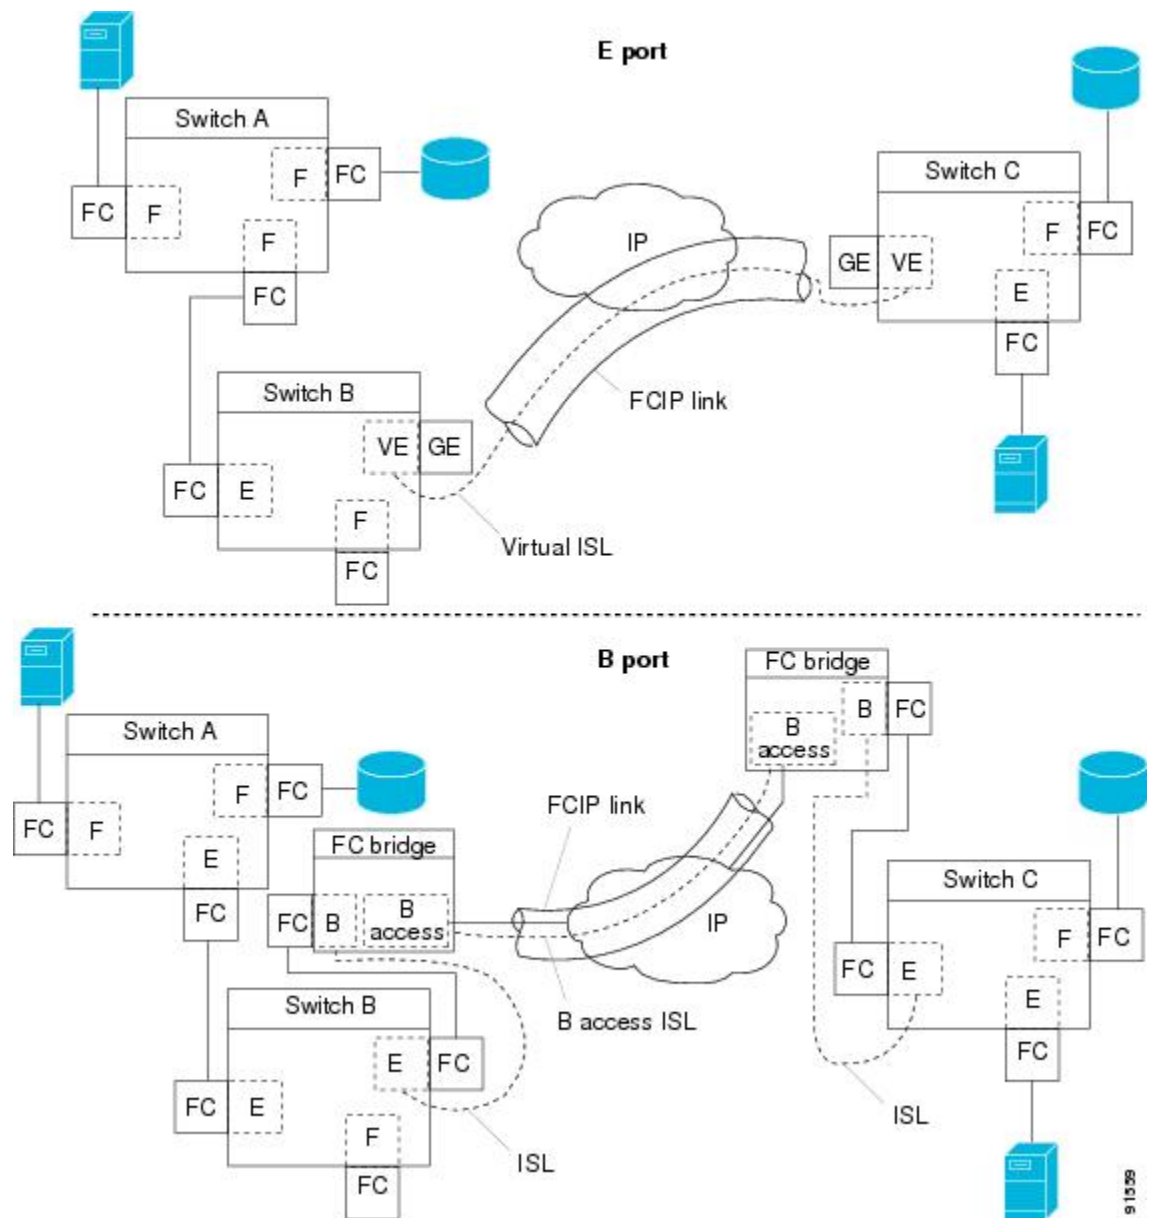
To establish an FCIP link with the peer, you can use the peer IP address option. This option configures both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.

To establish a peer connection, you must first create the FCIP interface and enter the `c onfig-if` submode.

FCIP B Port Interoperability Mode

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as Cisco's PA-FC-1G Fibre Channel port adapter and the SN 5428-2 storage router, implement a bridge port model to connect geographically dispersed fabrics. This model uses B port as described in the T11 Standard FC-BB-2.

[Figure 38-11](#) shows a typical SAN extension over an IP network.



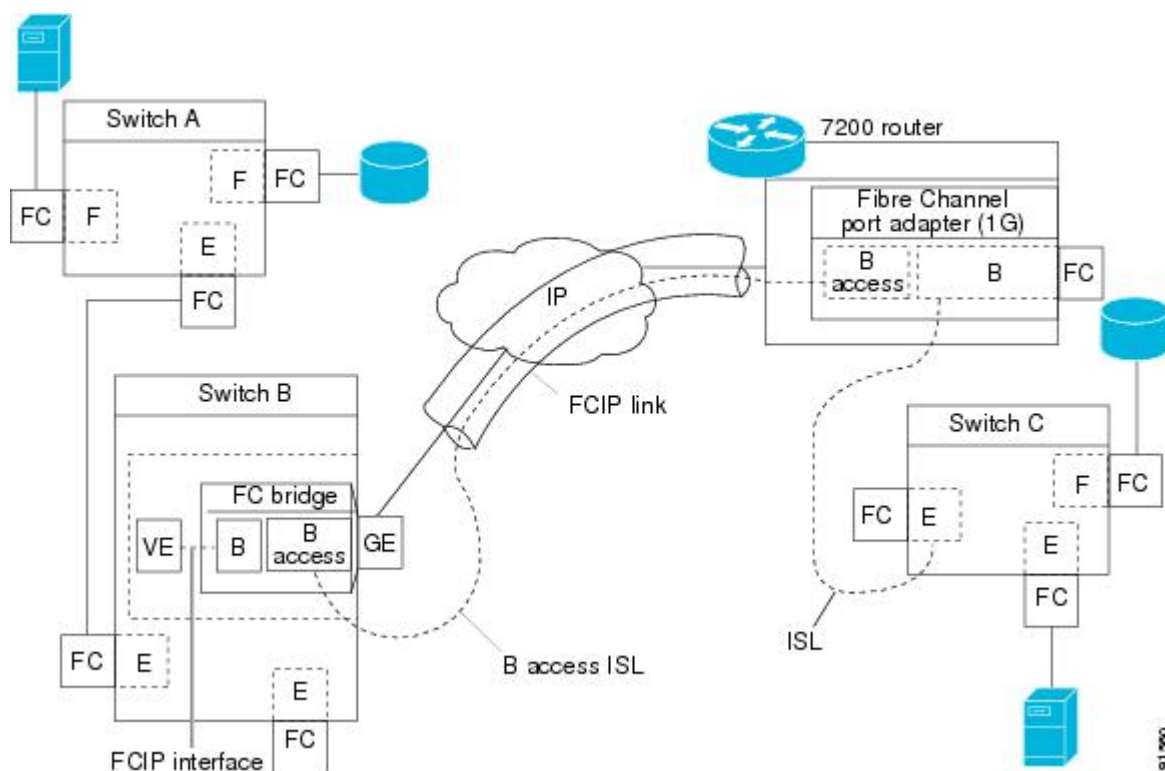
B ports bridge Fibre Channel traffic from a local E port to a remote E port without participating in fabric-related activities such as principal switch election, domain ID assignment, and Fibre Channel fabric shortest path first (FSPF) routing. For example, Class F traffic entering a SAN extender does not interact with the B port. The traffic is transparently propagated (bridged) over a WAN interface before exiting the remote B port. This bridge results in both E ports exchanging Class F information that ultimately leads to normal ISL behavior such as fabric merging and routing.

FCIP links between B port SAN extenders do not exchange the same information as FCIP links between E ports, and are therefore incompatible. This is reflected by the terminology used in FC-BB-2: *while VE ports establish a virtual ISL over an FCIP link, B ports use a B access ISL.*

The IPS module and MPS-14/2 module support FCIP links that originate from a B port SAN extender device by implementing the B access ISL protocol on a Gigabit Ethernet interface. Internally, the corresponding

virtual B port connects to a virtual E port that completes the end-to-end E port connectivity requirement (see [Figure 112: FCIP Link Terminating in a B Port Mode](#), on page 774).

Figure 112: FCIP Link Terminating in a B Port Mode



The B port feature in the IPS module and MPS-14/2 module allows remote B port SAN extenders to communicate directly with a Cisco MDS 9000 Family switch, eliminating the need for local bridge devices.

Quality of Service

The quality of service (QoS) parameter specifies the differentiated services code point (DSCP) value to mark all IP packets (type of service—TOS field in the IP header).

- The control DSCP value applies to all FCIP frames in the control TCP connection.
- The data DSCP value applies to all FCIP frames in the data connection.

If the FCIP link has only one TCP connection, that data DSCP value is applied to all packets in that connection.

E Ports

You can configure E ports in the same way you configure FCIP interfaces. The following features are also available for FCIP interfaces:

- An FCIP interface can be a member of any VSAN
- Trunk mode and trunk allowed VSANs
- PortChannels
- FSPF
- Fibre Channel domains (fcdomains)
- Importing and exporting the zone database from the adjacent switch

You can configure E ports in the same way you configure FCIP interfaces. The following features are also available for FCIP interfaces:

- An FCIP interface can be a member of any VSAN

See the Fabric Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Fabric Configuration Guide .

- Trunk mode and trunk allowed VSANs

See the Interfaces Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide.

- PortChannels
 - Multiple FCIP links can be bundled into a Fibre Channel PortChannel.
 - FCIP links and Fibre Channel links cannot be combined in one PortChannel. See the Security Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Security Configuration Guide.
- FSPF

See the Cisco MDS 9000 Family NX-OS Fabric Configuration Guide.

- Fibre Channel domains (fcdomains)

See the Cisco MDS 9000 Family NX-OS System Management Configuration Guide.

- Importing and exporting the zone database from the adjacent switch

See the Cisco MDS 9000 Family NX-OS System Management Configuration Guide.

FCIP Write Acceleration

The FCIP write acceleration feature enables you to significantly improve application write performance when storage traffic is routed over wide area networks using FCIP. When FCIP write acceleration is enabled, WAN throughput is maximized by minimizing the impact of WAN latency for write operations.



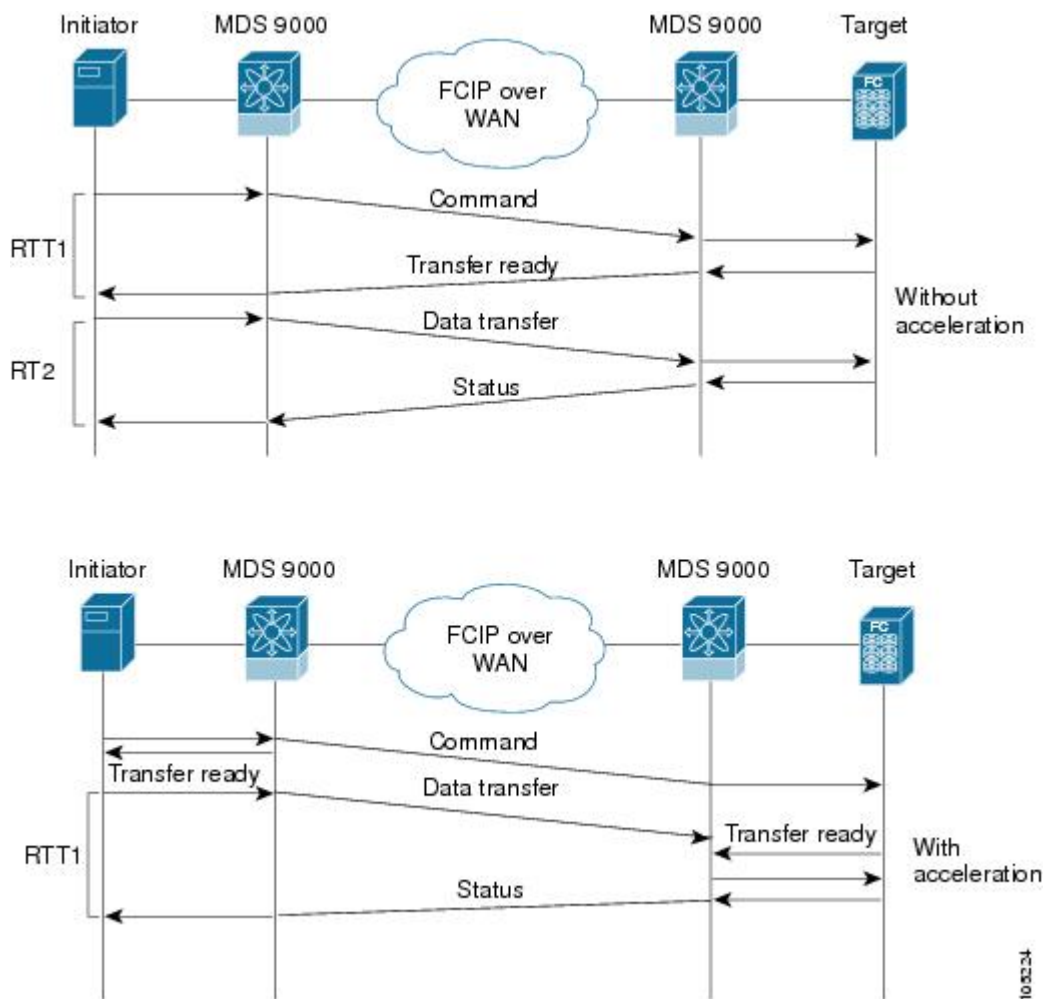
Note The write acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. If it is only enabled on one side of the FCIP tunnel the write acceleration feature will be turned operationally off.



Note IBM Peer to Peer Remote Copy (PPRC) is not supported with FCIP write acceleration.

The WRITE command (see [Figure 113: FCIP Link Write Acceleration, on page 776](#)), without write acceleration requires two round-trip transfers (RTT), while the WRITE command with write acceleration only requires one RTT. The maximum sized Transfer Ready is sent from the host side of the FCIP link back to the host before the WRITE command reaches the target. This enables the host to start sending the write data without waiting for the long latency over the FCIP link of the WRITE command and Transfer Ready. It also eliminates the delay caused by multiple Transfer Readys needed for the exchange going over the FCIP link.

Figure 113: FCIP Link Write Acceleration

**Tip**

FCIP write acceleration can be enabled for multiple FCIP tunnels if the tunnels are part of a dynamic PortChannel configured with channel mode active. FCIP write acceleration does not work if multiple non-PortChannel ISLs exist with equal weight between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or failed WRITE or READ operations.

**Tip**

Do not enable time stamp control on an FCIP interface with write acceleration configured.

**Note**

Write acceleration cannot be used across FSPF equal cost paths in FCIP deployments. Native Fibre Channel write acceleration can be used with PortChannels. Also, FCIP write acceleration can be used in PortChannels configured with channel mode active or constructed with PortChannel Protocol (PCP).

**Caution**

In Cisco MDS SAN-OS Release 2.0(1b) and later and NX-OS Release 4.x, FCIP write acceleration with FCIP ports as members of PortChannels are not compatible with the FCIP write acceleration in earlier releases.

FCIP Tape Acceleration

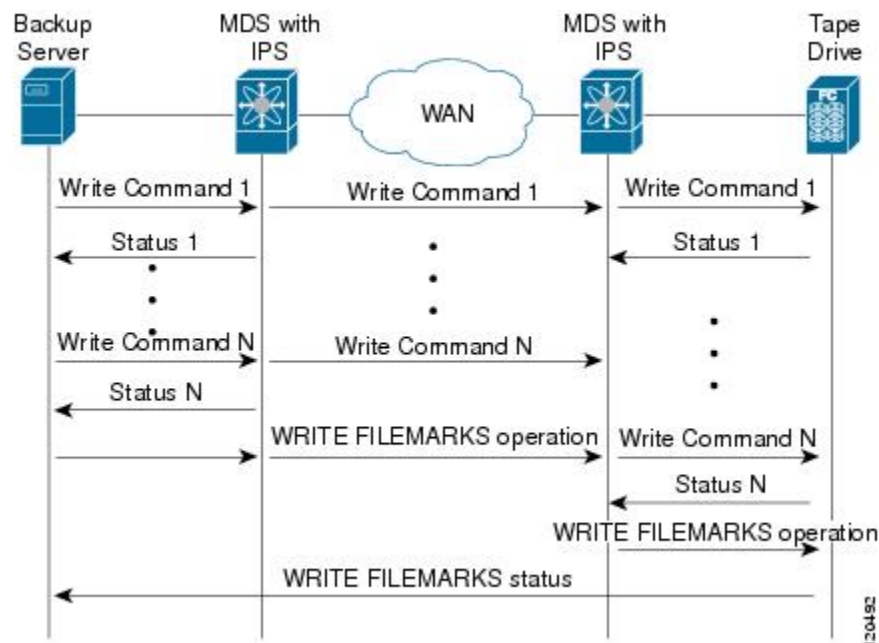
The FCIP write acceleration feature enables you to significantly improve application write performance when storage traffic is routed over wide area networks using FCIP. When FCIP write acceleration is enabled, WAN throughput is maximized by minimizing the impact of WAN latency for write operations. The write acceleration feature is disabled by default and must be enabled on both sides of the FCIP link.

Tapes are storage devices that store and retrieve user data sequentially. Cisco MDS NX-OS provides both tape write and read acceleration.

Applications that access tape drives normally have only one SCSI WRITE or READ operation outstanding to it. This single command process limits the benefit of the tape acceleration feature when using an FCIP tunnel over a long-distance WAN link. It impacts backup, restore, and restore performance because each SCSI WRITE or READ operation does not complete until the host receives a good status response from the tape drive. The FCIP tape acceleration feature helps solve this problem. It improves tape backup, archive, and restore operations by allowing faster data streaming between the host and tape drive over the WAN link.

In an example of tape acceleration for write operations, the backup server in (see [Figure 114: FCIP Link Tape Acceleration for Write Operations, on page 777](#)) issues write operations to a drive in the tape library. Acting as a proxy for the remote tape drives, the local Cisco MDS switch proxies a transfer ready to signal the host to start sending data. After receiving all the data, the local Cisco MDS switch proxies the successful completion of the SCSI WRITE operation. This response allows the host to start the next SCSI WRITE operation. This proxy method results in more data being sent over the FCIP tunnel in the same time period compared to the time taken to send data without proxying. The proxy method improves the performance on WAN links.

Figure 114: FCIP Link Tape Acceleration for Write Operations



At the tape end of the FCIP tunnel, another Cisco MDS switch buffers the command and data it has received. It then acts as a backup server to the tape drive by listening to a transfer ready from the tape drive before forwarding the data.

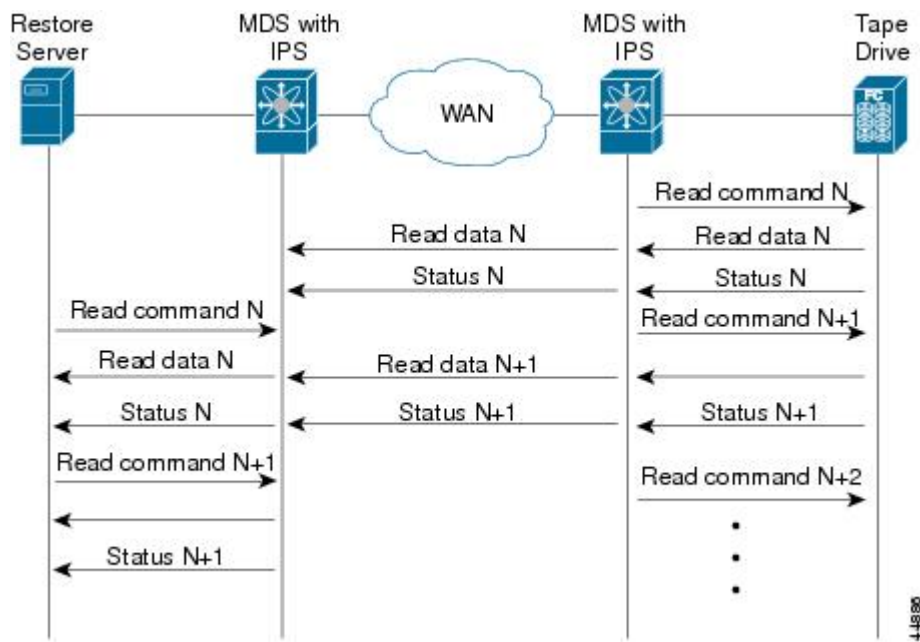
**Note**

In some cases such as a quick link up/down event (FCIP link, Server/Tape Port link) in a tape library environment that exports Control LUN or a Medium Changer as LUN 0 and tape drives as other LUNs, tape acceleration may not detect the tape sessions and may not accelerate these sessions. You need to keep the FCIP link disabled for a couple of minutes before enabling the link. This does not apply to tape environments where the tape drives are either direct FC attached or exported as LUN 0.

The Cisco NX-OS provides reliable data delivery to the remote tape drives using TCP/IP over the WAN. It maintains write data integrity by allowing the WRITE FILEMARKS operation to complete end-to-end without proxying. The WRITE FILEMARKS operation signals the synchronization of the buffer data with the tape library data. While tape media errors are returned to backup servers for error handling, tape busy errors are retried automatically by the Cisco NX-OS software.

In an example of tape acceleration for read operations, the restore server (see [Figure 115: FCIP Link Tape Acceleration for Read Operations, on page 778](#)) issues read operations to a drive in the tape library. During the restore process, the remote Cisco MDS switch at the tape end, in anticipation of more SCSI read operations from the host, sends out SCSI read operations on its own to the tape drive. The prefetched read data is cached at the local Cisco MDS switch. The local Cisco MDS switch on receiving SCSI read operations from the host, sends out the cached data. This method results in more data being sent over the FCIP tunnel in the same time period compared to the time taken to send data without read acceleration for tapes. This improves the performance for tape reads on WAN links.

Figure 115: FCIP Link Tape Acceleration for Read Operations



The Cisco NX-OS provides reliable data delivery to the restore application using TCP/IP over the WAN. While tape media errors during the read operation are returned to the restore server for error handling, the Cisco NX-OS software recovers from any other errors.



Note The FCIP tape acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. If it is only enabled on one side of the FCIP tunnel, the tape acceleration feature is turned operationally off.



Tip FCIP tape acceleration does not work if the FCIP port is part of a PortChannel or if there are multiple paths between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or broken write or read operations.



Caution When tape acceleration is enabled in an FCIP interface, a FICON VSAN cannot be enabled in that interface. Likewise, if an FCIP interface is up in a FICON VSAN, tape acceleration cannot be enabled on that interface.



Note When you enable the tape acceleration feature for an FCIP tunnel, the tunnel is reinitialized and the write and read acceleration feature is also automatically enabled.

In tape acceleration for writes, after a certain amount of data has been buffered at the remote Cisco MDS switch, the write operations from the host are flow controlled by the local Cisco MDS switch by not proxying the Transfer Ready. On completion of a write operation when some data buffers are freed, the local Cisco MDS switch resumes the proxying. Likewise, in tape acceleration for reads, after a certain amount of data has been buffered at the local Cisco MDS switch, the read operations to the tape drive are flow controlled by the remote Cisco MDS switch by not issuing any further reads. On completion of a read operation, when some data buffers are freed, the remote Cisco MDS switch resumes issuing reads.

The default flow control buffering uses the **automatic** option. This option takes the WAN latencies and the speed of the tape into account to provide optimum performance. You can also specify a flow control buffer size (the maximum buffer size is 12 MB).



Tip We recommend that you use the default option for flow-control buffering.



Tip Do not enable time-stamp control on an FCIP interface with tape acceleration configured.



Note If one end of the FCIP tunnel is running Cisco MDS SAN-OS Release 3.0(1) or later and NX-OS Release 4.x, and the other end is running Cisco MDS SAN-OS Release 2.x, and tape acceleration is enabled, then the FCIP tunnel will run only tape write acceleration, not tape-read acceleration.




Note In Cisco MDS NX-OS Release 4.2(1), the FCIP Tape Acceleration feature is not supported on FCIP back-to-back connectivity between MDS switches.

Tape Library LUN Mapping for FCIP Tape Acceleration

If a tape library provides logical unit (LU) mapping and FCIP tape acceleration is enabled, you must assign a unique LU number (LUN) to each physical tape drive accessible through a target port.

[Figure 116: FCIP LUN Mapping Example, on page 780](#) shows tape drives connected to Switch 2 through a single target port. If the tape library provides LUN mapping, then all the four tape drives should be assign unique LUNs.

Figure 116: FCIP LUN Mapping Example



For the mappings described in [Table 97: Correct LUN Mapping Example with Single Host Access, on page 780](#) and [Table 98: Incorrect LUN Mapping Example with Single Hosts Access, on page 781](#), Host 1 has access to Drive 1 and Drive 2, and Host 2 has access to Drive 3 and Drive 4.

[Table 97: Correct LUN Mapping Example with Single Host Access, on page 780](#) describes correct tape library LUN mapping.

Table 97: Correct LUN Mapping Example with Single Host Access

Host	LUN Mapping	Drive
Host 1	LUN 1	Drive 1
	LUN 2	Drive 2

Host	LUN Mapping	Drive
Host 2	LUN 3	Drive 3
	LUN 4	Drive 4

[Table 98: Incorrect LUN Mapping Example with Single Hosts Access, on page 781](#) describes incorrect tape library LUN mapping.

Table 98: Incorrect LUN Mapping Example with Single Hosts Access

Host	LUN Mapping	Drive
Host 1	LUN 1	Drive 1
	LUN 2	Drive 2
Host 2	LUN 1	Drive 3
	LUN 2	Drive 4

Another example setup is when a tape drive is shared by multiple hosts through a single tape port. For instance, Host 1 has access to Drive 1 and Drive 2, and Host 2 has access to Drive 2, Drive 3, and Drive 4. A correct LUN mapping configuration for such a setup is shown in [Table 99: Correct LUN Mapping Example with Multiple Host Access, on page 781](#).

Table 99: Correct LUN Mapping Example with Multiple Host Access

Host	LUN Mapping	Drive
Host 1	LUN 1	Drive 1
	LUN 2	Drive 2
Host 2	LUN 2	Drive 2
	LUN 3	Drive 3
	LUN 4	Drive 4

FCIP Compression

The FCIP compression feature allows IP packets to be compressed on the FCIP link if this feature is enabled on that link. By default the FCIP compression is disabled. When enabled, the software defaults to using the **auto** mode (if a mode is not specified).



Note The **auto** mode (default) selects the appropriate compression scheme based on the card type and bandwidth of the link (the bandwidth of the link configured in the FCIP profile's TCP parameters).

[Table 100: Algorithm Classification, on page 782](#) lists the modes used for different cards.

Table 100: Algorithm Classification

Mode	IPS Card	MPS 14/2 Card	MSM-18/4/MDS 9222i/SSN-16
mode1	SW	HW	HW
mode2	SW	SW	HW
mode3	SW	SW	HW



Note With SAN-OS Release 3.3(1) and later and NX-OS Release 4.x, all compression options on the MDS 9222i switch and the MSM-18/4 module, mean hardware compression. Starting with Release 4.2(1), only auto compression and mode 2 compression are supported on the MDS 9222i switch, the MSM-18/4 module, and the SSN-16 module.

Table 2-5 lists the performance settings for different cards.

Table 101: Performance Settings

Bandwidth	IPS Cards	MPS 14/2 Card	MSM-18/4/MDS 9222i/SSN-16
Any	-	-	auto
>25 Mbps	mode 1	mode 1	auto
10-25 Mbps	mode 2	mode 2	auto
10 Mbps	mode 3	mode 3	auto



Note The Cisco MDS 9216i and 9222i Switches also support the IP compression feature. The integrated supervisor module has the same hardware components that are available in the MPS-14/2 module.

**Caution**

The compression modes in Cisco SAN-OS Release 2.0(1b) and later and NX-OS Release 4.x are incompatible with the compression modes in Cisco SAN-OS Release 1.3(1) and earlier.

**Tip**

While upgrading from Cisco SAN-OS Release 1.x to Cisco SAN-OS Release 2.0(1b) or later and NX-OS Release 4.x, we recommend that you disable compression before the upgrade procedure, and then enable the required mode after the upgrade procedure.

If both ends of the FCIP link are running Cisco SAN-OS Release 2.0(1b) or later and NX-OS Release 4.x and you enable compression at one end of the FCIP tunnel, be sure to enable it at the other end of the link.

Default Settings

[Table 102: Default IKE Parameters](#), on [page 783](#) lists the default settings for IKE parameters.

Table 102: Default IKE Parameters

Parameters	Default
IKE	Disabled.
IKE version	IKE version 2.
IKE encryption algorithm	3DES.
IKE hash algorithm	SHA.
IKE authentication method	Not configurable (uses preshared Preshared keys).
IKE DH group identifier	Group 1.
IKE lifetime association	86,400 00 seconds (equals 24 hours).
IKE keepalive time for each peer (v2)	3,600 seconds (equals 1 hour).

[Table 103: Default IPsec Parameters](#), on [page 783](#) lists the default settings for IPsec parameters.

Table 103: Default IPsec Parameters

Parameters	Default
IPsec	Disabled.
Applying IPsec to the traffic.	Deny—allowing clear text.
IPsec PFS	Disabled.
IPsec global lifetime (traffic-volume)	450 Gigabytes.
IPsec global lifetime (time)	3,600 seconds (one hour).

Configuring FCIP

This section describes how to configure FCIP and includes the following topics:

Enabling FCIP

To begin configuring the FCIP feature, you must explicitly enable FCIP on the required switches in the fabric. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

The configuration and verification operations commands for the FCIP feature are only available when FCIP is enabled on a switch. When you disable this feature, all related configurations are automatically discarded.

To use the FCIP feature, you need to obtain the SAN extension over IP package license (SAN_EXTN_OVER_IP or SAN_EXTN_OVER_IP_IPS4) (see the Cisco Family NX-OS Licensing Guide). By default, the MDS 9222i and 9216i switches are shipped with the SAN extension over IP package license.



Note

In Cisco MDS SAN-OS Release 2.0 and later and NX-OS Release 4.x, there is an additional login prompt to log into a switch that is not a part of your existing fabric.

To create and manage FCIP links with DCNM-SAN, use the FCIP Wizard. Make sure that the IP Services Module is inserted in the required Cisco MDS 9000 Family switch, and that the Gigabit Ethernet interfaces on these switches are connected, and then verify the connectivity. The procedures for creating FCIP links using the FCIP Wizard are as follows:

- Select the endpoints.
- Choose the interfaces' IP addresses.
- Specify link attributes.
- (Optional) Enable FCIP write acceleration or FCIP compression.

To create FCIP links using the FCIP Wizard, follow these steps:

Procedure

Step 1

Click the FCIP Wizard icon in the DCNM-SAN toolbar (See [Figure 117: FCIP Wizard, on page 784](#)).

Figure 117: FCIP Wizard



Step 2

Choose the switches that act as endpoints for the FCIP link and click Next.

Step 3

Choose the Gigabit Ethernet ports on each switch that will form the FCIP link.

Step 4

If both Gigabit Ethernet ports are part of MPS-14/2 modules, check the Enforce IPSEC Security check box and set the IKE Auth Key. See the Security Configuration Guide, Cisco DCNM for SAN for information on IPsec and IKE.

Check the **Use Large MTU Size (Jumbo Frames)** option to use jumbo size frames of 2300. Since Fibre Channel frames are 2112, we recommended that you use this option. If you uncheck the box, the FCIP Wizard does not set the MTU size, and the default value of 1500 is set.

Note In Cisco MDS 9000 SAN-OS, Release 3.0(3), by default the **Use Large MTU Size (Jumbo Frames)** option is not selected.

- Step 5** Click Next.
You see the **IP Address/Route** input screen.
- Step 6** Select **Add IP Route** if you want to add an IP route, otherwise retain the defaults.
- Step 7** Click **Next**.
You see the TCP connection characteristics.
- Step 8** Set the minimum and maximum bandwidth settings and round-trip time for the TCP connections on this FCIP link.
You can measure the round-trip time between the Gigabit Ethernet endpoints by clicking the Measure button.
- Step 9** Check the Write Acceleration check box to enable FCIP write acceleration on this FCIP link.
See the [FCIP Write Acceleration, on page 775](#).
- Step 10** Check the Enable Optimum Compression check box to enable IP compression on this FCIP link.
See the [FCIP Compression, on page 782](#).
- Step 11** Check the Enable XRC Emulator check box to enable XRC emulator on this FCIP link.
For more information on XRC Emulator, see the Fabric Configuration Guide, Cisco DCNM for SAN.
- Step 12** Click Next.
- Step 13** Set the **Port VSAN** and click the **Trunk Mode** radio button for this FCIP link.
- Step 14** Click Finish to create this FCIP link.
-

Modifying an FCIP Link

Once you have created FCIP links using the FCIP wizard, you may need to modify parameters for these links. This procedure includes modifying the FCIP profiles as well as the FCIP link parameters. Each Gigabit Ethernet interface can have three active FCIP links at one time.

To modify an FCIP link, follow these steps on both switches:

Procedure

- Step 1** Configure the Gigabit Ethernet interface.
- Step 2** Create an FCIP profile, and then assign the Gigabit Ethernet interface's IP address to the profile.
- Step 3** Create an FCIP interface, and then assign the profile to the interface.
- Step 4** Configure the peer IP address for the FCIP interface.

- Step 5** Enable the interface.
-

Creating FCIP Profiles

To create an FCIP profile in switch 1, follow these steps:

Procedure

- Step 1** Verify that you are connected to a switch that contains an IPS module.
- Step 2** From DCNM-SAN, choose Switches > ISLs > FCIP in the Physical Attributes pane. From Device Manager, choose FCIP from the IP menu.
- Step 3** Click the Create Row button in DCNM-SAN or the Create button on Device Manager to add a new profile.
- Step 4** Enter the profile ID in the ProfileId field.
- Step 5** Enter the IP address of the interface to which you want to bind the profile.
- Step 6** Modify the optional TCP parameters, if desired. Refer to DCNM for SAN Online Help for explanations of these fields.
- Step 7** (Optional) Click the Tunnels tab and modify the remote IP address in the Remote IPAddress field for the endpoint to which you want to link.
- Step 8** Enter the optional parameters, if required.
- See the [FCIP Profiles, on page 768](#) for information on displaying FCIP profile information.
- Step 9** Click Apply Changes icon to save these changes.
-

Checking Trunk Status

By default, trunk mode is enabled in all Fibre Channel interfaces. However, trunk mode configuration takes effect only in E-port mode. You can configure trunk mode as on (enabled), off (disabled), or auto (automatic). The default trunk mode is on. The trunk mode configuration at the two ends of an ISL, between two switches, determines the trunking state of the link and the port modes at both ends.

To check the trunk status for the FCIP interface on Device Manager, follow these steps:

Procedure

- Step 1** Make sure you are connected to a switch that contains an IPS module.
- Step 2** Select FCIP from the IP menu.
- Step 3** Click the Trunk Config tab if it is not already selected. You see the FCIP Trunk Config dialog box. This shows the status of the interface.
- Step 4** Click the Trunk Failures tab if it is not already selected. You see the FCIP Trunk Failures dialog box.
-

Launching Cisco Transport Controller

Cisco Transport Controller (CTC) is a task-oriented tool used to install, provision, and maintain network elements. It is also used to troubleshoot and repair NE faults.

To launch CTC, follow these steps:

Procedure

-
- | | |
|---------------|------------------------------------------------------------|
| Step 1 | Right-click an ISL carrying optical traffic in the fabric. |
| Step 2 | Click Element Manager. |
| Step 3 | Enter the URL for the Cisco Transport Controller. |
| Step 4 | Click OK. |
-

Configuring TCP Parameters

You can control TCP behavior in a switch by configuring the TCP parameters that are described in this section.



Note When FCIP is sent over a WAN link, the default TCP settings may not be appropriate. In such cases, we recommend that you tune the FCIP WAN link by modifying the TCP parameters (specifically bandwidth, round-trip times, and CWM burst size).

This section includes the following topics:

Configuring Minimum Retransmit Timeout

You can control the minimum amount of time TCP waits before retransmitting. By default, this value is 200 milliseconds (msec).

Configuring Keepalive Timeout

You can configure the interval that the TCP connection uses to verify that the FCIP link is functioning. This ensures that an FCIP link failure is detected quickly even when there is no traffic.

If the TCP connection is idle for more than the specified time, then keepalive timeout packets are sent to ensure that the connection is active. This command can be used to tune the time taken to detect FCIP link failures.

You can configure the first interval during which the connection is idle (the default is 60 seconds). When the connection is idle for the configured interval, eight keepalive probes are sent at 1-second intervals. If no response is received for these eight probes and the connection remains idle throughout, that FCIP link is automatically closed.



Note Only the first interval (during which the connection is idle) can be changed.

Configuring Maximum Retransmissions

You can specify the maximum number of times a packet is retransmitted before TCP decides to close the connection.

Configuring Path MTUs

Path MTU (PMTU) is the minimum MTU on the IP network between the two endpoints of the FCIP link. PMTU discovery is a mechanism by which TCP learns of the PMTU dynamically and adjusts the maximum TCP segment accordingly (RFC 1191).

By default, PMTU discovery is enabled on all switches with a timeout of 3600 seconds. If TCP reduces the size of the maximum segment because of PMTU change, the reset-timeout specifies the time after which TCP tries the original MTU.

Configuring Selective Acknowledgments

TCP may experience poor performance when multiple packets are lost within one window. With the limited information available from cumulative acknowledgments, a TCP sender can only learn about a single lost packet per round trip. A selective acknowledgment (SACK) mechanism helps overcome the limitations of multiple lost packets during a TCP transmission.

The receiving TCP sends back SACK advertisements to the sender. The sender can then retransmit only the missing data segments. By default, SACK is enabled on Cisco MDS 9000 Family switches.

Configuring Window Management

The optimal TCP window size is automatically calculated using the maximum bandwidth parameter, the minimum available bandwidth parameter, and the dynamically measured round-trip time (RTT).



Note

The configured round-trip-time parameter determines the window scaling factor of the TCP connection. This parameter is only an approximation. The measured RTT value overrides the round trip time parameter for window management. If the configured round-trip-time is too small compared to the measured RTT, then the link may not be fully utilized due to the window scaling factor being too small.

The **min-available-bandwidth** parameter and the measured RTT together determine the threshold below which TCP aggressively maintains a window size sufficient to transmit at minimum available bandwidth.

The **max-bandwidth-mbps** parameter and the measured RTT together determine the maximum window size.



Note

Set the maximum bandwidth to match the worst-case bandwidth available on the physical link, considering other traffic that might be going across this link (for example, other FCIP tunnels, WAN limitations). Maximum bandwidth should be the total bandwidth minus all other traffic going across that link.

Configuring Monitoring Congestion

By enabling the congestion window monitoring (CWM) parameter, you allow TCP to monitor congestion after each idle period. The CWM parameter also determines the maximum burst size allowed after an idle period. By default, this parameter is enabled and the default burst size is 50 KB.

The interaction of bandwidth parameters and CWM and the resulting TCP behavior is outlined as follows:

- If the average rate of the Fibre Channel traffic over the preceding RTT is less than the min-available-bandwidth multiplied by the RTT, the entire burst is sent immediately at the min-available-bandwidth rate, provided no TCP drops occur.
- If the average rate of the Fibre Channel traffic is greater than min-available-bandwidth multiplied by the RTT, but less than max-bandwidth multiplied by the RTT, then if the Fibre Channel traffic is transmitted in burst sizes smaller than the configured CWM value the entire burst is sent immediately by FCIP at the max-bandwidth rate.
- If the average rate of the Fibre Channel traffic is larger than the min-available-bandwidth multiplied by the RTT and the burst size is greater than the CWM value, then only a part of the burst is sent immediately. The remainder is sent with the next RTT.

The software uses standard TCP rules to increase the window beyond the one required to maintain the min-available-bandwidth to reach the max-bandwidth.



Note The default burst size is 50 KB.



Tip We recommend that this feature remain enabled to realize optimal performance. Increasing the CWM burst size can result in more packet drops in the IP network, impacting TCP performance. Only if the IP network has sufficient buffering, try increasing the CWM burst size beyond the default to achieve lower transmit latency.

Configuring Estimating Maximum Jitter

Jitter is defined as a variation in the delay of received packets. At the sending side, packets are sent in a continuous stream with the packets spaced evenly apart. Due to network congestion, improper queuing, or configuration errors, this steady stream can become lumpy, or the delay between each packet can vary instead of remaining constant.

You can configure the maximum estimated jitter in microseconds by the packet sender. The estimated variation should not include network queuing delay. By default, this parameter is enabled in Cisco MDS switches when IPS modules or MPS-14/2 modules are present.

The default value is 1000 microseconds for FCIP interfaces.

Configuring Buffer Size

You can define the required additional buffering—beyond the normal send window size—that TCP allows before flow controlling the switch's egress path for the FCIP interface. The default FCIP buffer size is 0 KB.



Note Use the default if the FCIP traffic is passing through a high throughput WAN link. If you have a mismatch in speed between the Fibre Channel link and the WAN link, then time stamp errors occur in the DMA bridge. In such a situation, you can avoid time stamp errors by increasing the buffer size.

Assigning a Peer IP Address

The basic FCIP configuration uses the peer's IP address to configure the peer information. You can also specify the peer's port number to configure the peer information. If you do not specify a port, the default 3225 port number is used to establish connection. You can specify an IPv4 address or an IPv6 address.

Assign peer information based on the IPv4 address

To assign the peer information based on the IPv4 address and port number, follow these steps:

Procedure

-
- Step 1** Expand ISLs and select FCIP in the Physical Attributes pane.
You see the FCIP profiles and links in the Information pane.
From Device Manager, choose IP > FCIP.
You see the FCIP dialog box.
 - Step 2** Click the Tunnels tab. You see the FCIP link information.
 - Step 3** Click the Create Row icon in DCNM-SAN or the Create button in Device Manager.
You see the FCIP Tunnels dialog box.
 - Step 4** Set the ProfileID and TunnelID fields.
 - Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
 - Step 6** Check the PassiveMode check box if you do not want this end of the link to initiate a TCP connection.
 - Step 7** (Optional) Set the NumTCPCon field to the number of TCP connections from this FCIP link.
 - Step 8** (Optional) Check the Enable check box in the Time Stamp section and set the Tolerance field.
 - Step 9** (Optional) Set the other fields in this dialog box and click Create to create this FCIP link.
-

Assign peer information based on the IPv6 address

To assign the peer information based on the IPv6 address and port number, follow these steps:

Procedure

-
- Step 1** From DCNM-SAN, choose ISLs > FCIP from the Physical Attributes pane.
You see the FCIP profiles and links in the Information pane.
From Device Manager, choose IP > FCIP. You see the FCIP dialog box.
 - Step 2** Click the Tunnels tab. You see the FCIP link information.
 - Step 3** Click the Create Row icon in DCNM- SAN or the Create button in Device Manager.
You see the FCIP Tunnels dialog box.
 - Step 4** Set the ProfileID and TunnelID fields.
 - Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.

- Step 6** Check the PassiveMode check box if you do not want this end of the link to initiate a TCP connection.
- Step 7** (Optional) Set the NumTCPCon field to the number of TCP connections from this FCIP link.
- Step 8** (Optional) Check the Enable check box in the Time Stamp section and set the Tolerance field.
- Step 9** (Optional) Set the other fields in this dialog box and click Create to create this FCIP link.

Configuring Active Connections

You can configure the required mode for initiating a TCP connection. By default, the active mode is enabled to actively attempt an IP connection. If you enable the passive mode, the switch does not initiate a TCP connection but waits for the peer to connect to it. By default, the switch tries two TCP connections for each FCIP link.



Note Ensure that both ends of the FCIP link are not configured as passive mode. If both ends are configured as passive, the connection is not initiated.

Enabling Time Stamp Control

You can instruct the switch to discard packets that are outside the specified time. When enabled, this feature specifies the time range within which packets can be accepted. If the packet arrived within the range specified by this option, the packet is accepted. Otherwise, it is dropped.

By default, time stamp control is disabled in all switches in the Cisco MDS 9000 Family. If a packet arrives within a 2000 millisecond interval (+ or -2000 msec) from the network time, that packet is accepted.



Note The default value for packet acceptance is 2000 microseconds. If the **time-stamp** option is enabled, be sure to configure NTP on both switches (see the Cisco NX-OS Fundamentals Configuration Guide for more information).



Tip Do not enable time stamp control on an FCIP interface that has tape acceleration or write acceleration configured.

Configuring B Ports

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as Cisco's PA-FC-1G Fibre Channel port adapter and the SN 5428-2 storage router, implement a bridge port model to connect geographically dispersed fabrics. This model uses B port as described in the T11 Standard FC-BB-2. B ports bridge Fibre Channel traffic from a local E port to a remote E port without participating in fabric-related activities such as principal switch election, domain ID assignment, and Fibre Channel fabric shortest path first (FSPF) routing. The IPS module and MPS-14/2 module support FCIP links that originate from a B port SAN extender device by implementing the B access ISL protocol on a Gigabit Ethernet interface.

When an FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled.

To enable B port mode , follow these steps:

Procedure

-
- Step 1** Choose ISLs > FCIP from the Physical Attributes pane.
You see the FCIP profiles and links in the Information pane.
From Device Manager, choose IP > FCIP. You see the FCIP dialog box.
- Step 2** Click the Tunnels tab.
You see the FCIP link information.
- Step 3** Click the Create Row icon in DCNM-SAN or the Create button in Device Manager.
You see the FCIP Tunnels dialog box.
- Step 4** Set the ProfileID and TunnelID fields.
- Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
- Step 6** Check the PassiveMode check box if you do not want this end of the link to initiate a TCP connection.
- Step 7** (Optional) Set the NumTCPCon field to the number of TCP connections from this FCIP link.
- Step 8** Check the Enable check box in the B Port section of the dialog box and optionally check the KeepAlive check box if you want a response sent to an ELS Echo frame received from the FCIP peer.
- Step 9** (Optional) Set the other fields in this dialog box and click Create to create this FCIP link.
-

Configuring FCIP Write Acceleration

You can enable FCIP write acceleration when you create the FCIP link using the FCIP Wizard.

To enable write acceleration on an existing FCIP link, follow these steps:

Procedure

-
- Step 1** Choose ISLs > FCIP from the Physical Attributes pane on DCNM-SAN.
You see the FCIP profiles and links in the Information pane.
On Device Manager, choose IP > FCIP.
You see the FCIP dialog box.
- Step 2** Click the Tunnels (Advanced) tab.
You see the FICP link information.
- Step 3** Check or uncheck the Write Accelerator check box.
- Step 4** Choose the appropriate compression ratio from the IP Compression drop-down list.
- Step 5** Click the Apply Changes icon to save these changes.
-

Configuring FCIP Tape Acceleration

To enable FCIP tape acceleration, follow these steps:

Procedure

-
- Step 1** From DCNM-SAN, choose ISLs > FCIP from the Physical Attributes pane.
You see the FCIP profiles and links in the Information pane.
From Device Manager, choose IP > FCIP.
You see the FCIP dialog box.
- Step 2** Click the Tunnels tab. You see the FICP link information.
- Step 3** Click the Create Row icon in DCNM-SAN or the Create button in Device Manager.
You see the FCIP Tunnels dialog box.
- Step 4** Set the profile ID in the ProfileID field and the tunnel ID in the TunnelID fields.
- Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
- Step 6** Check the TapeAccelerator check box.
- Step 7** (Optional) Set the other fields in this dialog box and click Create to create this FCIP link.
-

Displaying FCIP Profile Information

To verify the FCIP interfaces and Extended Link Protocol (ELP) on Device Manager, follow these steps:

Procedure

-
- Step 1** Make sure you are connected to a switch that contains an IPS module.
- Step 2** Select FCIP from the Interface menu.
- Step 3** Click the Interfaces tab if it is not already selected. You see the FCIP Interfaces dialog box.
- Step 4** Click the ELP tab if it is not already selected. You see the FCIP ELP dialog box.
-

Field Descriptions for FCIP

This section describes the field description for FCIP.

FCIP Monitor

Field	Description
C3 Rx Bytes	The number of incoming bytes of data traffic.
C3 Tx Bytes	The number of outgoing bytes of data traffic.

Field	Description
CF Rx Bytes	The number of incoming bytes of control traffic.
CF Tx Bytes	The number of outgoing bytes of control traffic.
Rx Error	The number of inbound frames that contained errors preventing them from being deliverable to a higher-layer protocol.
Tx Error	The number of outbound frames that could not be transmitted because of errors.
RxDiscard	The number of inbound frames that were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol.
TxDiscard	The number of outbound frames that were chosen to be discarded even though no errors had been detected to prevent their being transmitted.

FCIP Interfaces

Field	Description
Description	Alias name for the interface as specified by a network manager.
PortVsan	The VSAN ID to which this interface is statically assigned.
Oper	The current operating mode of the port.
AutoChannelCreate	If checked, automatically create the PortChannel.
Admin	The desired state of the interface.
Oper	The current operational state of the interface.
FailureCause	The cause of current operational state of the port.
LastChange	When the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this value is N/A.
FICON Address	The FICON port address of this port.

FCIP Interfaces Trunk Failures

Field	Description
FailureCause	An entry is shown in this table if there is an error in the trunk status for the given VSAN.

FCIP FICON Configuration

Field	Description
Interface	This is a unique value that identifies the interface on this FCIP device to which this link pertains.
VSAN List Admin	This is the list of VSANs (in the range 1 through 2047) for which FICON tape acceleration is configured. Only VSANs with a cficonVsanEntry of CISCO-FICON-MIB present can be configured for FICON tape acceleration.
VSAN List Oper	This is the list of VSANs (in the range 1 through 2047) for which FICON tape acceleration is operationally ON.

FCIP Profiles

Field	Description
IP Address	The Internet address for this entity.
Port	A TCP port other than the FCIP well-known port on which the FCIP entity listens for new TCP connection requests.
SACK	Whether the TCP Selective Acknowledgement Option is enabled to allow the receiver end to acknowledge multiple lost frames in a single ACK, enabling faster recovery.
KeepAlive (s)	The TCP keep-alive timeout for all links within this entity.
ReTrans MinTimeout (ms)	The TCP minimum retransmit timeout for all the links on this entity.
ReTrans Max	The maximum number of times that the same item of data will be retransmitted over a TCP connection. If delivery is not acknowledged after this number of retransmissions then the connection is terminated.
Send BufSize (KB)	The aggregate TCP send window for all TCP connections on all Links within this entity. This value is used for egress flow control. When the aggregate of the data queued on all connections within this entity reaches this value, the sender is flow controlled.
Bandwidth Max (Kb)	This is an estimate of the bandwidth of the network pipe used for the B-D product computation, which lets us derive the TCP receive window to advertise.
Bandwidth Min (Kb)	The minimum available bandwidth for the TCP connections on the links within this entity.
Est Round Trip Time (us)	This is an estimate of the round trip delay of the network pipe used for the B-D product computation, which lets us derive the TCP receive window to advertise.
PMTU Enable	The path MTU discovery.
PMTU ResetTimeout (sec)	The time interval for which the discovered path MTU is valid, before MSS reverts back to the negotiated TCP value.

Field	Description
CWM Enable	If true, congestion window monitoring is enabled.
CWM BurstSize (KB)	The maximum burst sent after a TCP sender idle period.
Max Jitter	The maximum delay variation (not due to congestion) that can be experienced by TCP connections on this interface.

FCIP Tunnels

Field	Description
Interface	This identifies the interface on this FCIP device to which this link pertains.
Attached	The interface on which this FCIP link was initiated.
B Port Enable	If true, the B port mode is enabled on the local FCIP link.
B Port KeepAlive	If true, a message is sent in response to a (Fibre Channel) ELS Echo frame received from the peer. Some B Port implementations use ELS Echo request/response frames as link keep alive.
Remote IP Address	The Internet address for the remote FCIP entity.
Remote TCP Port	The remote TCP port to which the local FCIP entity will connect if and when it initiates a TCP connection setup for this link.
Spc Frames Enable	If true, the TCP active opener initiates FCIP special frames and the TCP passive opener responds to the FCIP special frames. If it is set to false, the FCIP special frames are neither generated nor responded to.
Spc Frames RemoteWWN	The world wide name of the remote FC fabric entity. If this is a zero length string then this link would accept connections from any remote entity. If a WWN is specified then this link would accept connections from a remote entity with this WWN.
Spc Frames Remote Profile Id	The remote FCIP entity's identifier.

FCIP Tunnels (Advanced)

Field	Description
Interface	The interface on which this FCIP link was initiated.
Timestamp Enable	If true, the timestamp in FCIP header is to be checked.
Timestamp Tolerance	The accepted time difference between the local time and the timestamp value received in the FCIP header. By default this value will be EDTOV/2. EDTOV is the Error_Detect_Timeout Value used for Fibre channel Ports as the timeout value for detecting an error condition.
Number Connections	The maximum number of TCP connections allowed on this link.

Field	Description
Passive	If false, this link endpoint actively tries to connect to the peer. If true, the link endpoint waits for the peer to connect to it.
QoS Control	The value to be set for the ToS field in IP header for the TCP control connection.
QoS Data	The value to be set for the ToS field in IP header for the TCP Data connection.
IP Compression	What algorithm is used, if any.
Write Accelerator	The write accelerator allows for enhancing SCSI write performance.
Tape Accelerator	If true, the tape accelerator (which allows for enhancing Tape write performance) is enabled.
Tape Accelerator Oper	Write acceleration is enabled for the FCIP link.
TapeRead Accelerator Oper	Enabled automatically when the tape accelerator Oper is active.
FlowCtrlBufSize Tape (KB)	The size of the flow control buffer (64 K to 32 MB). If set to 0, flow control buffer size is calculated automatically by the switch.
IPSec	Indicates whether the IP security has been turned on or off on this link.
XRC Emulator	Check to enable XRC emulator. It is disabled by default.
XRC Emulator Oper	Indicates the operational status of XRC emulator.

FCIP Tunnels (FICON TA)


Field	Description
Interface	A unique value that identifies the interface on this FCIP device to which this link pertains.
VSAN List Admin	The list of VSANs for which FICON tape acceleration is configured.
VSAN List Oper	The list of VSANs for which FICON tape acceleration is operationally on.

FCIP Tunnels Statistics

Field	Description
Interface	A unique value that identifies the interface on this FCIP device to which this link pertains.
Rx IPCompRatio	The IP compression ratio for received packets on the FCIP device. The value of this object will be presented as a floating point number with two digits after the decimal point.
Tx IPCompRatio	The IP compression ratio for transmitted packets on the FCIP device. The value of this object will be presented as a floating point number with two digits after the decimal point.

FCIP XRC Statistics

Field	Description
ProfileId	Unique ID of the profile.
Interface	Name of the interface.
RRSAccelerated	The number of read record set IUs accelerated.
RRSForwarded	Number of read record set IUs forwarded.
BusyStatus	Number of instances of busy status received from the control unit.
UnitCheckStatus	Number of instances of unit check status received from the control unit.
cfmFcipLinkExtXRCEStatsSelReset	Number of selective resets processed.
BufferAllocErrors	Number of buffer allocation errors.



CHAPTER 39

Configuring SAN Extension Tuner

- [Configuring the SAN Extension Tuner, on page 799](#)

Configuring the SAN Extension Tuner

The SAN Extension Tuner (SET) feature is unique to the Cisco MDS 9000 Family of switches. This feature helps you optimize FCIP performance by generating either direct access (magnetic disk) or sequential access (magnetic tape) SCSI I/O commands and directing such traffic to a specific virtual target. You can specify the size of the test I/O transfers and how many concurrent or serial I/Os to generate while testing. The SET reports the resulting I/Os per second (IOPS) and I/O latency, which helps you determine the number of concurrent I/Os needed to maximize FCIP throughput.

This chapter includes the following topics:

Information About the SAN Extension Tuner

The SAN extension tuner (SET) feature is unique to the Cisco MDS 9000 Family of switches. This feature helps you optimize FCIP performance by generating either direct access (magnetic disk) or sequential access (magnetic tape) SCSI I/O commands and directing such traffic to a specific virtual target. Applications such as remote copy and data backup use FCIP over an IP network to connect across geographically distributed SANs. SET is implemented in IPS ports. When enabled, this feature can be used to generate SCSI I/O commands (read and write) to the virtual target based on your configured options.



Note SAN Extension Tuner is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem, the Cisco Fabric Switch for IBM BladeCenter, and 16-Port Storage Services Node (SSN-16).




Note As of Cisco MDS SAN-OS Release 3.3(1a), SAN Extension Tuner is supported on the Multiservice Module (MSM) and the Multiservice Modular Switch.

Applications such as remote copy and data backup use FCIP over an IP network to connect across geographically distributed SANs. To achieve maximum throughput performance across the fabric, you can tune the following configuration parameters:

- The TCP parameters for the FCIP profile.
- The number of concurrent SCSI I/Os generated by the application.
- The transfer size used by the application over an FCIP link.

SET is implemented in IPS ports. When enabled, this feature can be used to generate SCSI I/O commands (read and write) to the virtual target based on your configured options (see [Figure 118: SCSI Command Generation to the Virtual Target](#), on page 800).

Figure 118: SCSI Command Generation to the Virtual Target



The SET feature assists with tuning by generating varying SCSI traffic workloads. It also measures throughput and response time per I/O over an FCIP link.

Before tuning the SAN fabric, be aware of the following guidelines:

- Following these implementation details:
 - The tuned configuration is not persistent.
 - The virtual N ports created do not register FC4 features supported with the name server. This is to avoid the hosts in the SAN from discovering these N ports as regular initiators or targets.
 - Login requests from other initiators in the SAN are rejected.
 - The virtual N ports do not implement the entire SCSI suite; it only implements the SCSI read and write commands.
 - Tuner initiators can only communicate with tuner targets.
- Verify that the Gigabit Ethernet interface is up at the physical layer (GBIC and Cable connected—an IP address is not required).
- Enable iSCSI on the switch (no other iSCSI configuration is required).
- Enable the interface (no other iSCSI interface configuration is required).

See [“Creating iSCSI Interfaces” section on page 40-36](#) for more information.

- Create an iSCSI interface on the Gigabit Ethernet interface and enable the interface (no other iSCSI interface configuration is required). See [“Creating iSCSI Interfaces” section on page 40-36](#) for more information.
- Configure the virtual N ports in a separate VSAN or zone as required by your network.


- Be aware that a separate VSAN with only virtual N ports is not required, but is recommended as some legacy HBAs may fail if logins to targets are rejected.
- Do not use same Gigabit Ethernet interface to configure virtual N ports and FCIP links—use different Gigabit Ethernet interfaces. While this is not a requirement, it is recommended as the traffic generated by the virtual N ports may interfere with the performance of the FCIP link.

This section includes the following topics:

SAN Extension Tuner Setup


[Figure 119: N Port Tuning Configuration Physical Example, on page 801](#) provides a sample physical setup in which the virtual N ports are created on ports that are not a part of the FCIP link for which the throughput and latency is measured.

Figure 119: N Port Tuning Configuration Physical Example



[Figure 120: Logical Example of N Port Tuning for a FCIP Link , on page 801](#) provides a sample logical setup in which the virtual N ports are created on ports that are not a part of the FCIP link for which the throughput and latency is measured.

Figure 120: Logical Example of N Port Tuning for a FCIP Link



Data Pattern

By default, an all-zero pattern is used as the pattern for data generated by the virtual N ports. You can optionally specify a file as the data pattern to be generated by selecting a data pattern file from one of three locations: the bootflash: directory, the volatile: directory, or the slot0: directory. This option is especially useful when testing compression over FCIP links. You can also use Canterbury corpus or artificial corpus files for benchmarking purposes.

Licensing Requirements for SAN Extension Tuner

To use the SET, you need to obtain the SAN_EXTN_OVER_IP license (see the Cisco Family NX-OS Licensing Guide).

The following table shows the licensing requirements for this feature:

License	License Description
SAN extension over IP package for IPS-8 modules <ul style="list-style-type: none"> • (SAN_EXTN_OVER_IP) SAN extension over IP package for IPS-4 modules <ul style="list-style-type: none"> • (SAN_EXTN_OVER_IP_IPS4) 	It comprises the SAN extension tuner features.
SAN extension over IP package for MPS-14/2 modules <ul style="list-style-type: none"> • (SAN_EXTN_OVER_IP_IPS2) 	This feature applies to the MPS-14/2 module and the fixed Cisco MDS 9216i Switch IP ports.
SAN extension over IP package for one MPS-18/4, one MPS-18/4 FIPS, or one SSN-16 engine in the Cisco MDS 9500 Series <ul style="list-style-type: none"> • (SAN_EXTN_OVER_IP_18_4) • (SAN_EXTN_OVER_IP_SSN16) 	This feature applies to the MPS-18/4, MPS-18/4 FIPS, or SSN-16 modules.

Default Settings

[Table 104: Default Tuning Parameters](#) , on page 802 lists the default settings for tuning parameters.

Table 104: Default Tuning Parameters

Parameters	Default
Tuning	Disabled
Transfer ready size	Same as the transfer size in the SCSI write command
Outstanding I/Os	1
Number of transactions	1
Data generation format	All-zero format

Configuring the SAN Extension Tuner

This section includes the following topics:

Tuning the FCIP Link

To tune the required FCIP link, follow these steps:

Procedure

-
- Step 1** Configure the nWWN for the virtual N ports on the switch.

- Step 2** Enable iSCSI on the interfaces on which you want to create the N ports.
- Step 3** Configure the virtual N ports on either side of the FCIP link.
- Step 4** Ensure that the virtual N ports are not visible to real initiators in the SAN. You can use zoning (see the Fabric Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Fabric Configuration Guide) to segregate the real initiators. Ensure that the zoning configuration is set up to allow the virtual N-ports to communicate with each other.
- Step 5** Start the SCSI read and write I/Os.
- Step 6** Add more N ports (as required) to other Gigabit Ethernet ports in the switch to obtain maximum throughput. One scenario that may require additional N ports is if you use FCIP PortChannels.
-

Using the SAN Extension Tuner Wizard

You can use the SAN Extension Tuner wizard to perform the these tasks:

- Configuring nWWN ports
- Enabling iSCSI
- Configuring Virtual N ports
- Assigning SCSI read and write CLI commands
- Assigning SCSI tape read and write CLI commands
- Configuring a data pattern for SCSI commands

To tune the required FCIP link using the SAN Extension Tuner Wizard in Cisco DCNM-SAN, follow these steps:

Procedure

- Step 1** Right-click a valid FCIP link in the Fabric pane, and then select **SAN Extension Tuner** from the drop-down list. You can also highlight the link and choose Tools > Other > SAN Extension Tuner.
- You see the Select Ethernet Port Pair dialog box.
- Step 2** Select the Ethernet port pairs that correspond to the FCIP link you want to tune and click Next.
- Note** The Ethernet ports you select should be listed as down.
- You see the Specify Parameters dialog box.
- Step 3** Create and activate a new zone to ensure that the virtual N ports are not visible to real initiators in the SAN by clicking Yes to the zone creation dialog box.
- Step 4** (Optional) Change the default settings for the transfer data size and the number of concurrent SCSI read and write commands as follows:
- a) Set Transfer Size to the number of bytes that you expect your applications to use over the FCIP link.
 - b) Set Read I/O to the number of concurrent SCSI read commands you expect your applications to generate over the FCIP link.
 - c) Set Write I/O to the number of concurrent outstanding SCSI write commands you expect your applications to generate over the FCIP link.
- Note** There is only one outstanding I/O at a time to the virtual N port that emulates the tape behavior.

- d) Check the Use Pattern File check box and select a file that you want to use to set the data pattern that is generated by the SAN extension tuner. See the [Data Pattern, on page 801](#).

Step 5 Click Next.

You see the Results dialog box.

Step 6 Click Start to start the tuner. The tuner sends a continuous stream of traffic until you click Stop.

Step 7 Click Show to see the latest tuning statistics. You can select this while the tuner is running or after you stop it.

Step 8 Click Stop to stop the SAN extension tuner.



CHAPTER 40

Configuring iSCSI

- [Configuring iSCSI, on page 805](#)

Configuring iSCSI

Cisco MDS 9000 Family IP storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch allows IP hosts to access Fibre Channel storage using the iSCSI protocol.



Note The iSCSI feature is specific to the IPS module and is available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors. The Cisco MDS 9216i switch and the 14/2 Multiprotocol Services (MPS-14/2) module also allow you to use Fibre Channel, FCIP, and iSCSI features. The MPS-14/2 module is available for use in any switch in the Cisco MDS 9200 Series or Cisco MDS 9500 Series.

This chapter includes the following topics:

Information About iSCSI

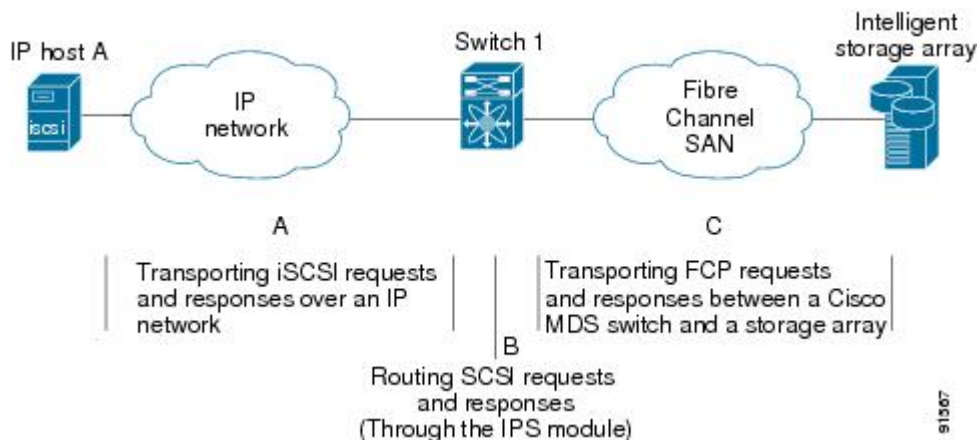
Cisco MDS 9000 Family IP Storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The iSCSI feature consists of routing iSCSI requests and responses between iSCSI hosts in an IP network and Fibre Channel storage devices in the Fibre Channel SAN that are accessible from any Fibre Channel interface of the Cisco MDS 9000 Family switch. Using the iSCSI protocol, the iSCSI driver allows an iSCSI host to transport SCSI requests and responses over an IP network. To use the iSCSI feature, you must explicitly enable iSCSI on the required switches in the fabric.



Note The iSCSI feature is not supported on the Cisco Fabric Switch for HP c-Class Bladesystem and Cisco Fabric Switch for IBM BladeCenter.

The iSCSI feature consists of routing iSCSI requests and responses between iSCSI hosts in an IP network and Fibre Channel storage devices in the Fibre Channel SAN that are accessible from any Fibre Channel interface of the Cisco MDS 9000 Family switch (see [Figure 121: Transporting iSCSI Requests and Responses for Transparent iSCSI Routing, on page 806](#)).

Figure 121: Transporting iSCSI Requests and Responses for Transparent iSCSI Routing




Each iSCSI host that requires access to storage through the IPS module or MPS-14/2 module needs to have a compatible iSCSI driver installed. Using the iSCSI protocol, the iSCSI driver allows an iSCSI host to transport SCSI requests and responses over an IP network. From the host operating system perspective, the iSCSI driver appears to be an SCSI transport driver similar to a Fibre Channel driver in the host.

The IPS module or MPS-14/2 module provides transparent SCSI routing. IP hosts using the iSCSI protocol can transparently access targets on the Fibre Channel network. It (see [Figure 121: Transporting iSCSI Requests and Responses for Transparent iSCSI Routing, on page 806](#)) provides an example of a typical configuration of iSCSI hosts connected to an IPS module or MPS-14/2 module through the IP network access Fibre Channel storage on the Fibre Channel SAN.

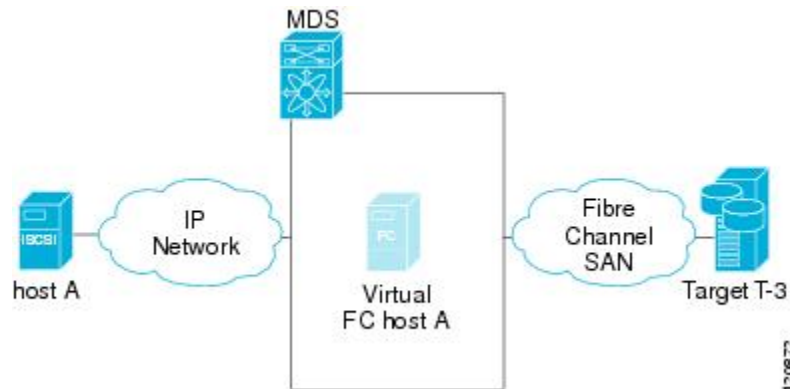
The IPS module or MPS-14/2 module create a separate iSCSI SAN view and Fibre Channel SAN view. For the iSCSI SAN view, the IPS module or MPS-14/2 module creates iSCSI virtual targets and then maps them to physical Fibre Channel targets available in the Fibre Channel SAN. They present the Fibre Channel targets to IP hosts as if the physical iSCSI targets were attached to the IP network (see [Figure 122: iSCSI SAN View—iSCSI Virtual Targets, on page 806](#)).

Figure 122: iSCSI SAN View—iSCSI Virtual Targets




For the Fibre Channel SAN view, the IPS module or MPS-14/2 module presents iSCSI hosts as a virtual Fibre Channel host. The storage devices communicate with the virtual Fibre Channel host similar to communications performed with real Fibre Channel hosts (see [Figure 123: Fibre Channel SAN View—iSCSI Host as an HBA, on page 807](#)).

Figure 123: Fibre Channel SAN View—iSCSI Host as an HBA



The IPS modules or MPS-14/2 modules transparently map the command between the iSCSI virtual target and the virtual Fibre Channel host (see Figure 124: iSCSI to FCP (Fibre Channel) Routing, on page 807).

Figure 124: iSCSI to FCP (Fibre Channel) Routing



Routing SCSI from the IP host to the Fibre Channel storage device consists of the following main actions:

- The iSCSI requests and responses are transported over an IP network between the hosts and the IPS module or MPS-14/2 module.
- The SCSI requests and responses are routed between the hosts on an IP network and the Fibre Channel storage device (converting iSCSI to FCP and vice versa). The IPS module or MPS-14/2 module performs this conversion and routing.
- The FCP requests or responses are transported between the IPS module or MPS-14/2 module and the Fibre Channel storage devices.



Note FCP (the Fibre Channel equivalent of iSCSI) carries SCSI commands over a Fibre Channel SAN. Refer to the IETF standards for IP storage at <http://www.ietf.org> for information on the iSCSI protocol.

About iSCSI Configuration Limits

iSCSI configuration has the following limits:

- The maximum number of iSCSI and iSLB initiators supported in a fabric is 2000.
- The maximum number of iSCSI and iSLB initiators supported is 200 per port.

- The maximum number of iSCSI and iSLB sessions supported by an IPS port in either transparent or proxy initiator mode is 500.
- The maximum number of iSCSI and iSLB session support by switch is 5000.
- The maximum number of iSCSI and iSLB targets supported in a fabric is 6000.

Presenting Fibre Channel Targets as iSCSI Targets

The IPS module or MPS-14/2 module presents physical Fibre Channel targets as iSCSI virtual targets, allowing them to be accessed by iSCSI hosts. The module presents these targets in one of the two ways:

- **Dynamic mapping**—Automatically maps all the Fibre Channel target devices/ports as iSCSI devices. Use this mapping to create automatic iSCSI target names.
- **Static mapping**—Manually creates iSCSI target devices and maps them to the whole Fibre Channel target port or a subset of Fibre Channel LUNs. With this mapping, you must specify unique iSCSI target names.

Static mapping should be used when iSCSI hosts should be restricted to subsets of LUs in the Fibre Channel targets and/or iSCSI access control is needed (see the [iSCSI Access Control, on page 814](#)). Also, static mapping allows the configuration of transparent failover if the LUs of the Fibre Channel targets are reachable by redundant Fibre Channel ports (see the [Transparent Target Failover, on page 825](#)).



Note

The IPS module or MPS-14/2 module does not import Fibre Channel targets to iSCSI by default. Either dynamic or static mapping must be configured before the IPS module or MPS-14/2 module makes Fibre Channel targets available to iSCSI initiators.

Dynamic Mapping

When you configure dynamic mapping the IPS module or MPS-14/2 module imports all Fibre Channel targets to the iSCSI domain and maps each physical Fibre Channel target port as one iSCSI target. That is, all LUs accessible through the physical storage target port are available as iSCSI LUs with the same LU number (LUN) as in the physical Fibre Channel target port.

The iSCSI target node name is created automatically using the iSCSI qualified name (IQN) format. The iSCSI qualified name is restricted to a maximum name length of 223 alphanumeric characters and a minimum length of 16 characters.

The IPS module or MPS-14/2 module creates an IQN formatted iSCSI target node name using the following conventions because the name must be unique in the SAN:

- IPS Gigabit Ethernet ports that are not part of a Virtual Router Redundancy Protocol (VRRP) group or PortChannel use this format:

```
iqn.1987-05.com.cisco:05.<mgmt-ip-address>.<slot#>-<port#>-<sub-intf#>.<Target-pWWN>
```

- IPS ports that are part of a VRRP group use this format:

```
iqn.1987-05.com.cisco:05.vrrp-<vrrp-ID#>-<vrrp-IP-addr>.<Target-pWWN>
```

- Ports that are part of a PortChannel use this format:

```
iqn.1987-02.com.cisco:02.<mgmt-ip-address>.pc-<port-ch-sub-intf#>.<Target-pWWN>
```




Note If you have configured a switch name, then the switch name is used instead of the management IP address. If you have not configured a switch name, the management IP address is used.

With this convention, each IPS port in a Cisco MDS 9000 Family switch creates a unique iSCSI target node name for the same Fibre Channel target port in the SAN.

For example, if an iSCSI target was created for a Fibre Channel target port with pWWN 31:00:11:22:33:44:55:66 and that pWWN contains LUN 0, LUN 1, and LUN 2, those LUNs would become available to an IP host through the iSCSI target node name `iqn.1987-05.com.cisco:05.MDS_switch_management_IP_address.01-01.3100112233445566` (see [Figure 125: Dynamic Target Mapping](#), on page 809).

Figure 125: Dynamic Target Mapping




Note Each iSCSI initiator may not have access to all targets depending on the configured access control mechanisms (see the [iSCSI Access Control](#), on page 814).

Static Mapping

You can manually (statically) create an iSCSI target by assigning a user-defined unique iSCSI node name to it. The iSCSI qualified name is restricted to a minimum length of 16 characters and a maximum of 223 characters. A statically mapped iSCSI target can either map the whole Fibre Channel target port (all LUNs in the target port mapped to the iSCSI target), or it can contain one or more LUs from a Fibre Channel target port (see [Figure 126: Statically Mapped iSCSI Targets](#), on page 810).

Figure 126: Statically Mapped iSCSI Targets



Presenting iSCSI Hosts as Virtual Fibre Channel Hosts

The IPS module or MPS-14/2 module connects to the Fibre Channel storage devices on behalf of the iSCSI host to send commands and transfer data to and from the storage devices. These modules use a virtual Fibre Channel N port to access the Fibre Channel storage devices on behalf of the iSCSI host. iSCSI hosts are identified by either iSCSI qualified name (IQN) or IP address.

Initiator Identification

iSCSI hosts can be identified by the IPS module or MPS-14/2 module using the following:

- iSCSI qualified name (IQN)

An iSCSI initiator is identified based on the iSCSI node name it provides in the iSCSI login. This mode can be useful if an iSCSI host has multiple IP addresses and you want to provide the same service independent of the IP address used by the host. An initiator with multiple IP addresses (multiple network interface cards—NICs) has one virtual N port on each IPS port to which it logs in.

- IP address

An iSCSI initiator is identified based on the IP address of the iSCSI host. This mode is useful if an iSCSI host has multiple IP addresses and you want to provide different service-based on the IP address used by the host. It is also easier to get the IP address of a host compared to getting the iSCSI node name. A virtual N port is created for each IP address it uses to log in to iSCSI targets. If the host using one IP address logs in to multiple IPS ports, each IPS port will create one virtual N port for that IP address.

Initiator Presentation Modes

Two modes are available to present iSCSI hosts in the Fibre Channel fabric: transparent initiator mode and proxy initiator mode.

- In transparent initiator mode, each iSCSI host is presented as one virtual Fibre Channel host. The benefit of transparent mode is it allows a finer level of Fibre Channel access control configuration (similar to managing a “real” Fibre Channel host). Because of the one-to-one mapping from iSCSI to Fibre Channel, each host can have different zoning or LUN access control on the Fibre Channel storage device.
- In proxy initiator mode, there is only one virtual Fibre Channel host per one IPS port and all iSCSI hosts use that to access Fibre Channel targets. In a scenario where the Fibre Channel storage device requires

explicit LUN access control for every host, the static configuration for each iSCSI initiator can be overwhelming. In this case, using the proxy initiator mode simplifies the configuration.


Caution

Enabling proxy initiator mode of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the [Changing iSCSI Interface Parameters and the Impact on Load Balancing, on page 823](#).

The Cisco MDS switches support the following iSCSI session limits:

- The maximum number of iSCSI sessions on a switch is 5000.
- The maximum number of iSCSI sessions per IPS port in transparent initiator mode is 500.
- The maximum number of iSCSI sessions per IPS port in proxy initiator mode is 500.
- The maximum number of concurrent sessions an IPS port can create is five (but the total number of sessions that can be supported is 500).


Note


If more than five iSCSI sessions try to come up simultaneously on a port, the initiator receives a temporary error and later retries to create a session.

Transparent Initiator Mode

Each iSCSI host is presented as one virtual Fibre Channel host (that is, one Fibre Channel N port). The benefit of transparent mode is it allows a finer-level of Fibre Channel access control configuration. Because of the one-to-one mapping from iSCSI to Fibre Channel, each host can have different zoning or LUN access control on the Fibre Channel storage device.

When an iSCSI host connects to the IPS module or MPS-14/2 module, a virtual host N port (HBA port) is created for the host (see [Figure 127: Virtual Host HBA Port, on page 811](#)). Every Fibre Channel N port requires a unique Node WWN and Port WWN.

Figure 127: Virtual Host HBA Port



After the virtual N port is created with the WWNs, a fabric login (FLOGI) is done through the virtual iSCSI interface of the IPS port. After the FLOGI is completed, the virtual N port is online in the Fibre Channel SAN and virtual N port is registered in the Fibre Channel name server. The IPS module or MPS-14/2 module registers the following entries in the Fibre Channel name server:

- IP address of the iSCSI host in the IP-address field on the name server
- IQN of the iSCSI host in the symbolic-node-name field of the name server
- SCSI_FCP in the FC-4 type field of the name server
- Initiator flag in the FC-4 feature of the name server
- Vendor-specific iSCSI GW flag in the FC-4 type field to identify the N-port device as an iSCSI gateway device in the name server.

When all the iSCSI sessions from the iSCSI host are terminated, the IPS modules or MPS-14/2 modules perform an explicit Fabric logout (FLOGO) to remove the virtual N-port device from the Fibre Channel SAN (this indirectly de-registers the device from the Fibre Channel name server).

For every iSCSI session from the host to the iSCSI virtual target there is a corresponding Fibre Channel session to the real Fibre Channel target. There are three iSCSI hosts (see [Figure 127: Virtual Host HBA Port, on page 811](#)), and all three of them connect to the same Fibre Channel target. There is one Fibre Channel session from each of the three virtual Fibre Channel hosts to the target.

WWN Assignment for iSCSI Initiators

An iSCSI host is mapped to an N port's WWNs by one of the following mechanisms:

- Dynamic mapping (default)
- Static mapping

Dynamic Mapping

With dynamic mapping, an iSCSI host is mapped to a dynamically generated port WWN (pWWN) and node WWN (nWWN). Each time the iSCSI host connects it might be mapped to a different WWN. Use this option if no access control is required on the Fibre Channel target device (because the target device access control is usually configured using the host WWN).

The WWNs are allocated from the MDS switch's WWN pool. The WWN mapping to the iSCSI host is maintained as long as the iSCSI host has at least one iSCSI session to the IPS port. When all iSCSI sessions from the host are terminated and the IPS module or MPS-14/2 module performs an FLOGO for the virtual N port of the host, the WWNs are released back to the switch's Fibre Channel WWN pool. These addresses are then available for assignment to other iSCSI hosts requiring access to the Fibre Channel Fabric.

The following are three dynamic initiator modes are supported:

- iSCSI—Dynamic initiators are treated as iSCSI initiators and can access dynamic virtual targets and configured iSCSI virtual targets.
- iSLB—Dynamic initiators are treated as iSLB initiators.
- Deny —Dynamic initiators are not allowed to log in to the MDS switch.

iSCSI dynamic mapping is the default mode of operation. This configuration is distributed using CFS.



Note

Configuring dynamic initiator modes is supported only through the CLI, not through Device Manager or Cisco DCNM for SAN.

Static Mapping

With static mapping, an iSCSI host is mapped to a specific pWWN and nWWN. This mapping is maintained in persistent storage and each time the iSCSI host connects, the same WWN mapping is used. This mode is required if you use access control on the target device.

You can implement static mapping in one of two ways:

- User assignment—You can specify your own unique WWN by providing them during the configuration process.
- System assignment—You can request that the switch provide a WWN from the switch's Fibre Channel WWN pool and keep the mapping in its configuration.




Tip We recommend using the **system-assign** option. If you manually assign a WWN, you must ensure its uniqueness. You should not use any previously assigned WWNs.

Proxy Initiator Mode

In the event that the Fibre Channel storage device requires explicit LUN access control for every host use the transparent initiator mode (presenting one iSCSI host as one Fibre Channel host). Every iSCSI host has to be configured statically. This can mean several configuration tasks for each iSCSI host. If you do not need explicit LUN access control, using the proxy initiator mode simplifies the configuration.

In this mode, only one virtual host N port (HBA port) is created per IPS port. All the iSCSI hosts connecting to that IPS port will be multiplexed using the same virtual host N port (see [Figure 128: Multiplexing IPS Ports, on page 813](#)). This mode simplifies the task of statically binding WWNs. LUN mapping and assignment on the Fibre Channel storage array must be configured to allow access from the proxy virtual N port's pWWN for all LUNs used by each iSCSI initiator that connects through this IPS port. The LUN is then assigned to each iSCSI initiator by configuring iSCSI virtual targets (see the [Static Mapping, on page 809](#)) with LUN mapping and iSCSI access control (see the [iSCSI Access Control, on page 814](#)).

Figure 128: Multiplexing IPS Ports



Proxy initiator mode can be configured on a per IPS port basis, in which case only iSCSI initiators terminating on that IPS port will be in this mode.

When an IPS port is configured in proxy-initiator mode, fabric login (FLOGI) is done through the virtual iSCSI interface of the IPS port. After the FLOGI is completed, the proxy-initiator virtual N port is online in the Fibre Channel fabric and virtual N port is registered in the Fibre Channel name server. The IPS module or MPS-14/2 module registers the following entries in the Fibre Channel name server:

- iSCSI interface name iSCSI slot /port is registered in the symbolic-node-name field of the name server
- SCSI_FCP in the FC-4 type field of the name server
- Initiator flag in the FC-4 feature of the name server
- Vendor specific flag (iscsi-gw) in the FC-4 type field to identify the N-port device as an iSCSI gateway device in the name server

Similar to transparent initiator mode, the user can provide a pWWN and nWWN or request a system assigned WWN for the proxy initiator N port.


Caution

Enabling the proxy initiator mode of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the [Changing iSCSI Interface Parameters and the Impact on Load Balancing, on page 823](#).

VSAN Membership for iSCSI

VSAN membership can be configured for an iSCSI interface, called the port VSAN. All the iSCSI devices that connect to this interface automatically become members of this VSAN, if it is not explicitly configured in a VSAN. The default port VSAN of an iSCSI interface is VSAN 1. Similar to Fibre Channel devices, iSCSI devices have two mechanisms by which VSAN membership can be defined.

- iSCSI host—VSAN membership to iSCSI host. (This method takes precedent over the iSCSI interface).
- iSCSI interface—VSAN membership to iSCSI interface. (All iSCSI hosts connecting to this iSCSI interface inherit the interface VSAN membership if the host is not configured in any VSAN by the iSCSI host method).

Advanced VSAN Membership for iSCSI Hosts

An iSCSI host can be a member of multiple VSANs. In this case, multiple virtual Fibre Channel hosts are created, one in each VSAN in which the iSCSI host is a member. This configuration is useful when certain resources such as Fibre Channel tape devices need to be shared among different VSANs.

iSCSI Access Control

Two methods of access control are available for iSCSI devices. Depending on the initiator mode used to present the iSCSI hosts in the Fibre Channel fabric, either or both of the access control methods can be used.

- Fiber Channel zoning-based access control—Fibre Channel zoning has been extended to support iSCSI devices, and this extension has the advantage of having a uniform, flexible access control mechanism across the whole SAN. In the case of iSCSI, multiple iSCSI devices may be connected behind an iSCSI interface. Interface-based zoning may not be useful because all iSCSI devices behind the interface will automatically be within the same zone.
- iSCSI ACL-based access control—iSCSI-based access control is applicable only if static iSCSI virtual targets are created. For a static iSCSI target, you can configure a list of iSCSI initiators that are allowed to access the targets. By default, static iSCSI virtual targets are not accessible to any iSCSI host.

Depending on the initiator mode used to present the iSCSI hosts in the Fibre Channel fabric, either or both the access control mechanisms can be used.

The following topics are included in this section:

Fibre Channel Zoning-Based Access Control

Cisco SAN-OS Release 3.x and NX-OS Release 4.1(1b) VSAN and zoning concepts have been extended to cover both Fibre Channel devices and iSCSI devices. Zoning is the standard access control mechanism for Fibre Channel devices, which is applied within the context of a VSAN. Fibre Channel zoning has been extended to support iSCSI devices, and this extension has the advantage of having a uniform, flexible access control mechanism across the whole SAN.

Common mechanisms for identifying members of a Fibre Channel zone are the following:

- Fibre Channel device pWWN.
- Interface and switch WWN. Device connecting via that interface is within the zone.

In the case of iSCSI, multiple iSCSI devices may be connected behind an iSCSI interface. Interface-based zoning may not be useful because all the iSCSI devices behind the interface will automatically be within the same zone.

In transparent initiator mode (where one Fibre Channel virtual N port is created for each iSCSI host as described in the [Transparent Initiator Mode, on page 811](#)), if an iSCSI host has static WWN mapping then the standard Fibre Channel device pWWN-based zoning membership mechanism can be used.

Zoning membership mechanism has been enhanced to add iSCSI devices to zones based on the following:

- IPv4 address/subnet mask
- IPv6 address/prefix length
- iSCSI qualified name (IQN)
- Symbolic-node-name (IQN)

For iSCSI hosts that do not have a static WWN mapping, the feature allows the IP address or iSCSI node name to be specified as zone members. Note that iSCSI hosts that have static WWN mapping can also use these features. IP address based zone membership allows multiple devices to be specified in one command by providing the subnet mask.



Note In proxy initiator mode, all iSCSI devices connecting to an IPS port gain access to the Fibre Channel fabric through a single virtual Fibre Channel N port. Zoning based on the iSCSI node name or IP address will not have any effect. If zoning based on pWWN is used, then all iSCSI devices connecting to that IPS port will be put in the same zone. To implement individual initiator access control in proxy initiator mode, configure an iSCSI ACL on the virtual target (see the [iSCSI-Based Access Control, on page 815](#)).

iSCSI-Based Access Control

iSCSI-based access control is applicable only if static iSCSI virtual targets are created (see the [Static Mapping, on page 809](#)). For a static iSCSI target, you can configure a list of iSCSI initiators that are allowed to access the targets.

By default, static iSCSI virtual targets are not accessible to any iSCSI host. You must explicitly configure accessibility to allow an iSCSI virtual target to be accessed by all hosts. The initiator access list can contain one or more initiators. The iSCSI initiator can be identified by one of the following mechanisms:

- iSCSI node name
- IPv4 address and subnet
- IPv6 address

**Note**

For a transparent mode iSCSI initiator, if both Fibre Channel zoning and iSCSI ACLs are used, then for every static iSCSI target that is accessible to the iSCSI host, the initiator's virtual N port should be in the same Fibre Channel zone as the Fibre Channel target.

Enforcing Access Control

IPS modules and MPS-14/2 modules use both iSCSI and Fibre Channel zoning-based access control lists to enforce access control. Access control is enforced both during the iSCSI discovery phase and the iSCSI session creation phase. Access control enforcement is not required during the I/O phase because the IPS module or MPS-14/2 module is responsible for the routing of iSCSI traffic to Fibre Channel.

If the iSCSI target is a static mapped target, the IPS module or MPS-14/2 module verifies if the iSCSI host is allowed within the access list of the iSCSI target. If the IP host does not have access, its login is rejected. If the iSCSI host is allowed, it validates if the virtual Fibre Channel N port used by the iSCSI host and the Fibre Channel target mapped to the static iSCSI virtual target are in the same Fibre Channel zone.

If the iSCSI target is an autogenerated iSCSI target, then the IPS module or MPS-14/2 module extracts the WWN of the Fibre Channel target from the iSCSI target name and verifies if the initiator and the Fibre Channel target is in the same Fibre Channel zone or not. If they are, then access is allowed.

The IPS module or MPS-14/2 module uses the Fibre Channel virtual N port of the iSCSI host and does a zone-enforced name server query for the Fibre Channel target WWN. If the FC ID is returned by the name server, then the iSCSI session is accepted. Otherwise, the login request is rejected.

iSCSI Session Authentication

The IPS module or MPS-14/2 module supports the iSCSI authentication mechanism to authenticate the iSCSI hosts that request access to the storage devices. By default, the IPS modules or MPS-14/2 modules allow CHAP or None authentication of iSCSI initiators. If authentication is always used, you must configure the switch to allow only CHAP authentication.

For CHAP user name or secret validation, you can use any method supported and allowed by the Cisco MDS AAA infrastructure. AAA authentication supports a RADIUS, TACACS+, or local authentication device.

The **aaa authentication iscsicommand** enables AAA authentication for the iSCSI host and specifies the method to use. See Cisco MDS 9000 Family NX-OS Security Configuration Guide.

iSCSI Immediate Data and Unsolicited Data Features

Cisco MDS switches support the iSCSI immediate data and unsolicited data features if requested by the initiator during the login negotiation phase. Immediate data is iSCSI write data contained in the data segment of an iSCSI command protocol data unit (PDU), such as combining the write command and write data together in one PDU. Unsolicited data is iSCSI write data that an initiator sends to the iSCSI target, such as an MDS switch, in an iSCSI data-out PDU without having to receive an explicit ready to transfer (R2T) PDU from the target.

These two features help reduce I/O time for small write commands because it removes one round-trip between the initiator and the target for the R2T PDU. As an iSCSI target, the MDS switch allows up to 64 KB of unsolicited data per command. This is controlled by the FirstBurstLength parameter during iSCSI login negotiation phase.

If an iSCSI initiator supports immediate data and unsolicited data features, these features are automatically enabled on the MDS switch with no configuration required.

Cisco MDS switches support the following advanced features for iSCSI interfaces:

iSCSI Listener Port

You can configure the TCP port number for the iSCSI interface that listens for new TCP connections. The default port number is 3260. Once you change the TCP port number, the iSCSI port only accepts TCP connections on the newly configured port.

TCP Tuning Parameters

You can configure the following TCP parameters:

- Minimum retransmit timeout (See the [“Configuring Minimum Retransmit Timeout”](#) section on page 38-26 for more information).
- Keepalive timeout.
- Maximum retransmissions (See the [“Configuring Maximum Retransmissions”](#) section on page 38-27 for more information).
- Path MTU (See the [“Configuring Path MTUs”](#) section on page 38-27 for more information).
- SACK (SACK is enabled by default for iSCSI TCP configurations).
- Window management (The iSCSI defaults are max-bandwidth is 1 Gbps, min-available-bandwidth is 70 Mbps, and round-trip-time is 1 msec). (See the [“Configuring Window Management”](#) section on page 38-28 for more information).
- Buffer size (The iSCSI default send buffer size is 4096 KB) (See the [“Configuring Buffer Size”](#) section on page 38-30 for more information).
- Window congestion monitoring (enabled by default and the default burst size is 50 KB) (See the [“Configuring Monitoring Congestion”](#) section on page 38-29 for more information).
- Maximum delay jitter (enabled by default and the default time is 500 microseconds).

iSCSI Routing Modes

Cisco MDS 9000 Family switches support multiple iSCSI routing modes. Each mode negotiates different operational parameters, has different advantages and disadvantages, and is suitable for different usages.

- Pass-thru mode

In pass-thru mode, the port on the IPS module or MPS 14/2 module converts and forwards read data frames from the Fibre Channel target to the iSCSI host frame-by-frame without buffering. This means that one data-in frame received is immediately sent out as one iSCSI data-in PDU.

In the opposite direction, the port on the IPS module or MPS 14/2 module limits the maximum size of iSCSI write data-out PDU that the iSCSI host can send to the maximum data size that the Fibre Channel target specifies that it can receive. The result is one iSCSI data-out PDU received sent out as one Fibre Channel data frame to the Fibre Channel target.

The absence of buffering in both directions leads to an advantage of lower forwarding latency. However, a small maximum data segment length usually results in lower data transfer performance from the host because of a higher processing overhead by the host system. Another benefit of this mode is iSCSI data digest can be enabled. This helps protect the integrity of iSCSI data carried in the PDU over what TCP checksum offers.

- Store-and-forward mode (default)

In store-and-forward mode, the port on the IPS module or MPS 14/2 module assembles all the Fibre Channel data frames of an exchange to build one large iSCSI data-in PDU before forwarding it to the iSCSI client.

In the opposite direction, the port on the IPS module or MPS 14/2 module does not impose a small data segment size on the host so the iSCSI host can send an iSCSI data-out PDU of any size (up to 256 KB). The port then waits until the whole iSCSI data-out PDU is received before it converts, or splits, the PDU, and forwards Fibre Channel frames to the Fibre Channel target.

The advantage of this mode is higher data transfer performance from the host. The disadvantages are higher transfer latency and that the iSCSI data digest (CRC) cannot be used.



Note The store-and-forward mode is the default forwarding mode.

- Cut-through mode

Cut-through mode improves the read operation performance over store-and-forward mode. The port on the IPS module or MPS 14/2 module achieves this by forwarding each Fibre Channel data-in frame to the iSCSI host as it is received without waiting for the whole exchange complete. There is no difference for write data-out operations from store-and-forward mode.

Figure 129: iSCSI Routing Modes, on page 818 compares the messages exchanged by the iSCSI routing modes.

Figure 129: iSCSI Routing Modes

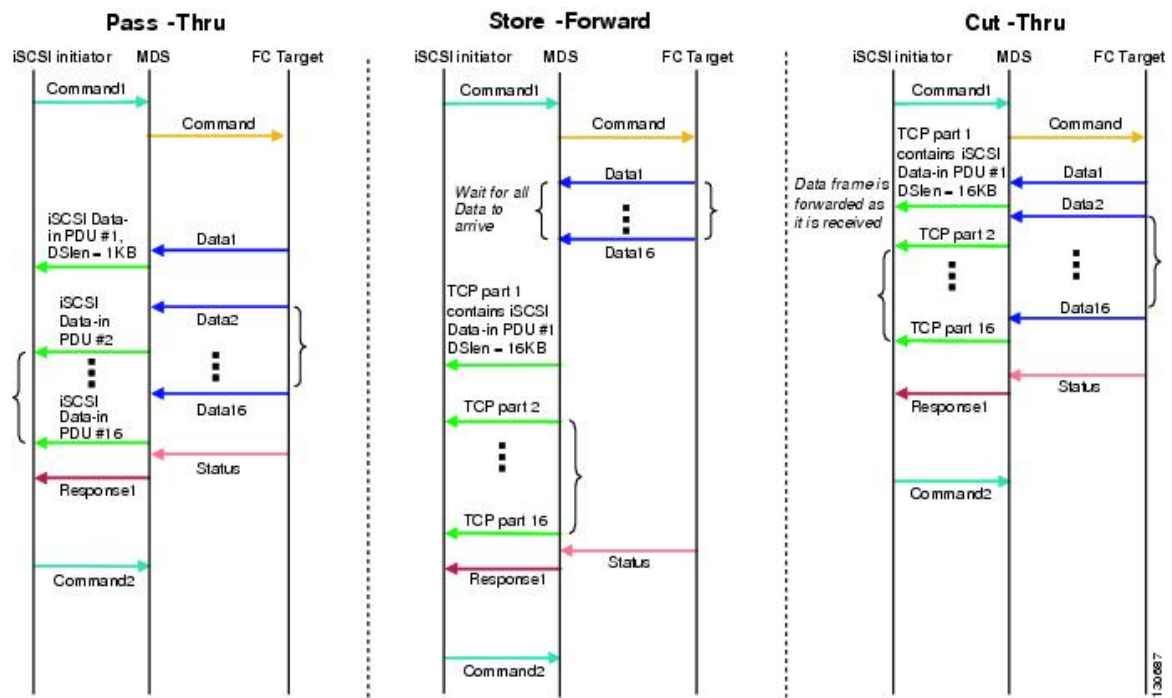


Table 105: Comparison of iSCSI Routing Modes, on page 819 compares the advantages and disadvantages of the different iSCSI routing modes.

Table 105: Comparison of iSCSI Routing Modes

Mode	Advantages	Disadvantages
Pass-thru	Low-latency Data digest can be used	Lower data transfer performance.
Store-and-forward	Higher data transfer performance	Data digest cannot be used.
Cut-thru	Improved read performance over store-and-forward	If the Fibre Channel target sent read data for different commands interchangeably, data of the first command is forwarded in cut-thru mode but the data of subsequent commands is buffered and the behavior is the same as store-and-forward mode. Data digest cannot be used.

**Caution**

Changing the forwarding mode of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the [Changing iSCSI Interface Parameters and the Impact on Load Balancing, on page 823](#).

About iSLB

The iSCSI server load balancing (iSLB) feature provides a means to easily configure large scale iSCSI deployments containing hundreds or even thousands of initiators. iSLB provides the following features:

- The iSLB initiator configuration is simplified with support for initiator targets and auto-zones.
- Cisco Fabric Services (CFS) eliminates the need for manual configuration by distributing the iSLB initiator configuration among all MDS switches in the fabric.
- Dynamic load balancing of iSLB initiators is available using iSCSI login redirect and VRRP.

When not using iSLB, configuring iSCSI requires the following:

- You need to perform multiple configuration steps on the MDS switch, including the following:
 - Initiator configuration using static pWWN and VSAN.
 - Zoning configuration for initiators and targets.
 - Optional create virtual target and give access to the initiator.
 - Configuration of target LUN mapping and masking on the storage system for the initiator based on the static pWWN created for the initiator on the MDS switch.
- You need to duplicate the configuration manually on multiple MDS switches.
- There is no load balancing for IPS ports. For example:
 - The Virtual Router Redundancy Protocol (VRRP) only supports active and backup, not load balancing.
 - You must use multiple VRRP groups and configure hosts in different groups.

iSLB provides the following features:

- The iSLB initiator configuration is simplified with support for initiator targets and auto-zones.

- Cisco Fabric Services (CFS) eliminates the need for manual configuration by distributing the iSLB initiator configuration among all MDS switches in the fabric.

**Note**

Only statically mapped iSLB initiator configuration is distributed throughout the fabric using CFS. Dynamically and statically mapped iSCSI initiator configurations are not distributed.

- Dynamic load balancing of iSLB initiators is available using iSCSI login redirect and VRRP.

About iSLB Initiators

iSLB initiators provide the following features in addition to those supported by iSCSI initiators:

- An iSLB initiator also supports iSLB virtual targets.
- Initiator targets—These targets are configured for a particular initiator.
- Load balancing using iSCSI login redirect and VRRP—If iSCSI login redirect is enabled, the IPS Manager redirects incoming sessions to the best interface based on the calculated load for each interface.
- Configuration distribution to other switches using CFS.

iSLB initiators provide the following features in addition to those supported by iSCSI initiators:

- An iSLB initiator also supports iSLB virtual targets. These targets are very similar to iSCSI virtual targets with the exception that they do not include the advertise interface option and as a result are distributable using CFS.
- Initiator targets—These targets are configured for a particular initiator.
- Load balancing using iSCSI login redirect and VRRP—If load balancing is enabled, the IPS Manager redirects incoming sessions to the best interface based on the calculated load for each interface.
- Configuration distribution to other switches using CFS.

Assigning WWNs to iSLB Initiators

An iSLB host is mapped to an N port's WWNs by one of the following mechanisms:

- Dynamic mapping (default)
- Static mapping

**Note**

Assigning WWNs for iSLB initiators is the same as for iSCSI initiators. For information on dynamic and static mapping, see the [WWN Assignment for iSCSI Initiators, on page 812](#).

**Tip**

We recommend using the **SystemAssign system-assign** option. If you manually assign a WWN, you must ensure its uniqueness. You should not use any previously assigned WWNs.

See the [Configuring iSLB Using Device Manager, on page 848](#).

iSLB Initiator Targets

You can configure initiator targets using the device alias or the pWWN. You can also optionally specify one or more of the following optional parameters:

- Secondary pWWN
- Secondary device alias
- LUN mapping
- IQN
- VSAN identifier



Note The VSAN identifier is optional if the target is online. If the target is not online, the VSAN identifier is required.

In addition, you can disable auto-zoning.

If you configure an IQN for an initiator target, then that name is used to identify the initiator target. Otherwise, a unique IQN is generated for the initiator target.

iSLB Session Authentication

The IPS module and MPS-14/2 module support the iSLB authentication mechanism to authenticate iSLB hosts that request access to storage. By default, the IPS module and MPS-14/2 module allow CHAP or None authentication of iSCSI initiators. If authentication is always used, you must configure the switch to allow only CHAP authentication.

For CHAP user name or secret validation you can use any method supported and allowed by the Cisco MDS AAA infrastructure. AAA authentication supports RADIUS, TACACS+, or a local authentication device.



Note Specifying the iSLB session authentication is the same as for iSCSI. See the [iSCSI Session Authentication, on page 816](#).


About Load Balancing Using VRRP

You can configure Virtual Router Redundancy Protocol (VRRP) load balancing for iSLB. The host is configured with a VRRP address as the portal address. When the VRRP master port receives the first iSCSI session from an initiator, it assigns a backup port to serve that particular host. The information is synchronized to all switches through CFS if recovery is needed when a master port fails. The initiator gets a temporary redirect iSCSI login response. The host then logs in to the backup port at its physical IP address. All iSCSI interfaces in a VRRP group that has load balancing enabled must have the same interface VSAN, authentication, proxy initiator mode, and forwarding mode.

You can configure Virtual Router Redundancy Protocol (VRRP) load balancing for iSLB.

[Figure 130: iSLB Initiator Load Balancing Example, on page 822](#) shows an example of load balancing using iSLB.

Figure 130: iSLB Initiator Load Balancing Example



The host is configured with a VRRP address as the portal address. When the VRRP master port receives the first iSCSI session from an initiator, it assigns a backup port to serve that particular host. This information is synchronized to all switches through CFS if recovery is needed when a master port fails. The initiator gets a temporary redirect iSCSI login response. The host then logs in to the backup port at its physical IP address. If the backup port goes down, the host will revert to the master port. The master port knows through CFS that the backup port has gone down and redirects the host to another backup port.



Note If an Ethernet PortChannel is configured between the IPS module and an Ethernet switch, the load balancing policy on the Ethernet switch must be based on source/destination IP address only, not port numbers, for load balancing with VRRP to operate correctly.



Note An initiator can also be redirected to the physical IP address of the master interface.



Tip iSLB VRRP load balancing is based on the number of iSLB initiators and not number of sessions. Any iSLB initiator that has more targets configured than the other iSLB initiators (resulting in more sessions) should be configured with a higher load metric. For example, you can increase the load metric of the iSLB initiator with more targets to 3000 from the default value of 1000.

**Caution**

A Gigabit Ethernet interface configured for iSLB can only be in one VRRP group because redirected sessions do not carry information about the VRRP IP address or group. This restriction allows the slave backup port to uniquely identify the VRRP group to which it belongs.

Changing iSCSI Interface Parameters and the Impact on Load Balancing

All iSCSI interfaces in a VRRP group that has load balancing enabled must have the same interface VSAN, authentication, proxy initiator mode, and forwarding mode. When you need to change any of these parameters for the iSCSI interfaces in a VRRP group, you must do so one interface at a time. During the transition time when the parameter is changed on some interfaces in the VRRP group and not the others, the master port does not redirect new initiators and instead handles them locally.

**Caution**

Changing the VSAN, proxy initiator, authentication, and forwarding mode for iSCSI interfaces in a VRRP group can cause sessions to go down multiple times.

VRRP Load Balancing Algorithm For Selecting Gigabit Ethernet Interfaces

When the VRRP master receives an iSCSI session request from an initiator, it first checks for an existing mapping to one of the interfaces in that VRRP group. If such a mapping exists, the VRRP master redirects the initiator to that interface. If no such mapping exists, the VRRP master selects the least loaded interface and updates the selected interface's load with the initiator's iSLB metric (weight).

**Note**

The VRRP master interface is treated specially and it needs to take a lower load compared to the other interfaces. This is to account for the redirection work performed by the master interface for every session. A new initiator is assigned to the master interface only if the following is true for every other interface: $\text{VRRP backup interface load} > [2 * \text{VRRP master interface load} + 1]$

About iSLB Configuration Distribution Using CFS

You can distribute the configuration for iSLB initiators and initiator targets on an MDS switch. This feature lets you synchronize the iSLB configuration across the fabric from the console of a single MDS switch. The iSCSI initiator idle timeout, global authentication, and iSCSI dynamic initiator mode parameters are also distributed. CFS distribution is disabled by default.

Configuration for iSLB initiators and initiator targets on an MDS switch can be distributed using the Cisco Fabric Services (CFS). This feature allows you to synchronize the iSLB configuration across the fabric from the console of a single MDS switch. The iSCSI initiator idle timeout, iSCSI dynamic initiator mode, and global authentication parameters are also distributed. CFS distribution is disabled by default.

After enabling the distribution, the first configuration starts an implicit session. All server configuration changes entered thereafter are stored in a temporary database and applied to all switches in the fabric (including the originating one) when you explicitly commit the database.

When CFS is enabled for iSLB, the first iSLB configuration operation starts a CFS session and locks the iSLB configuration in the fabric. The configuration changes are applied to the pending configuration database.

When you make the changes to the fabric, the pending configuration is distributed to all the switches in the fabric. Each switch then validates the configuration. This check ensures the following:

- The VSANs assigned to the iSLB initiators are configured on all the switches.
- The static WWNs configured for the iSLB initiators are unique and available on all the switches.
- The iSLB initiator node names do not conflict with the iSCSI initiators on all the switches.

After the check completes successfully, all the switches commit the pending configuration to the running configuration. If any check fails, the entire commit fails.



Note iSLB is only fully supported when CFS is enabled. Using iSLB auto-zoning without enabling CFS mode may cause traffic disruption when any zone set is activated.



Note CFS does not distribute non-iSLB initiator configurations or import Fibre Channel target settings.

Non-iSLB virtual targets will continue to support advertised interfaces option.



Tip The pending changes are only available in the volatile directory and are discarded if the switch is restarted.

Locking the Fabric

The first action that modifies the existing configuration creates the pending configuration and locks the feature in the fabric. Once you lock the fabric, the following conditions apply:

- No other user can make any configuration changes to this feature.
- A pending configuration is created by copying the active configuration. Modifications from this point on are made to the pending configuration and remain there until you commit the changes to the active configuration (and other switches in the fabric) or discard them.



Note iSCSI configuration changes are not allowed when an iSLB CFS session is active.

CFS Merge Process

When two fabrics merge, CFS attempts to merge the iSLB configuration from both the fabrics. A designated switch (called the *dominant switch*) in one fabric sends its iSLB configuration to a designated switch (called the *subordinate switch*) in the other fabric. The subordinate switch compares its running configuration to the received configuration for any conflicts. If no conflicts are detected, it merges the two configurations and sends it to all the switches in both the fabrics. Each switch then validates the configuration. This check ensures the following:

- VSANs assigned to the iSLB initiators are configured on all the switches.
- The static WWNs configured for the iSLB initiators are unique and available on all the switches.
- The iSLB initiator node names have no conflicts with iSCSI initiators on all the switches.

If this check completes successfully, the subordinate switch directs all the switches to commit the merged configuration to running configuration. If any check fails, the merge fails.

iSLB CFS Merge Status Conflicts

Merge conflicts may occur. User intervention is required for the following merge conflicts:

- The iSCSI global authentication or iSCSI initiator idle timeout parameters are not configured the same in the two fabrics.
- The same iSLB initiator is configured differently in the two fabrics.
- An iSLB initiator in one fabric has the same name as an iSCSI initiator in the other fabric.
- Duplicate pWWN/nWWN configuration is detected in the two fabric. For example, a pWWN/nWWN configured for an iSLB initiator on one fabric is configured for an iSCSI initiator or a different iSLB initiator in the other fabric.
- A VSAN configured for an iSLB initiator in one fabric does not exist in the other fabric.



Tip Check the syslog for details on merge conflicts.

User intervention is not required when the same iSLB initiator has a different set of non-conflicting initiator targets. The merged configuration is the union of all the initiator targets.

iSCSI High Availability

The following high availability features are available for iSCSI configurations:

Transparent Target Failover

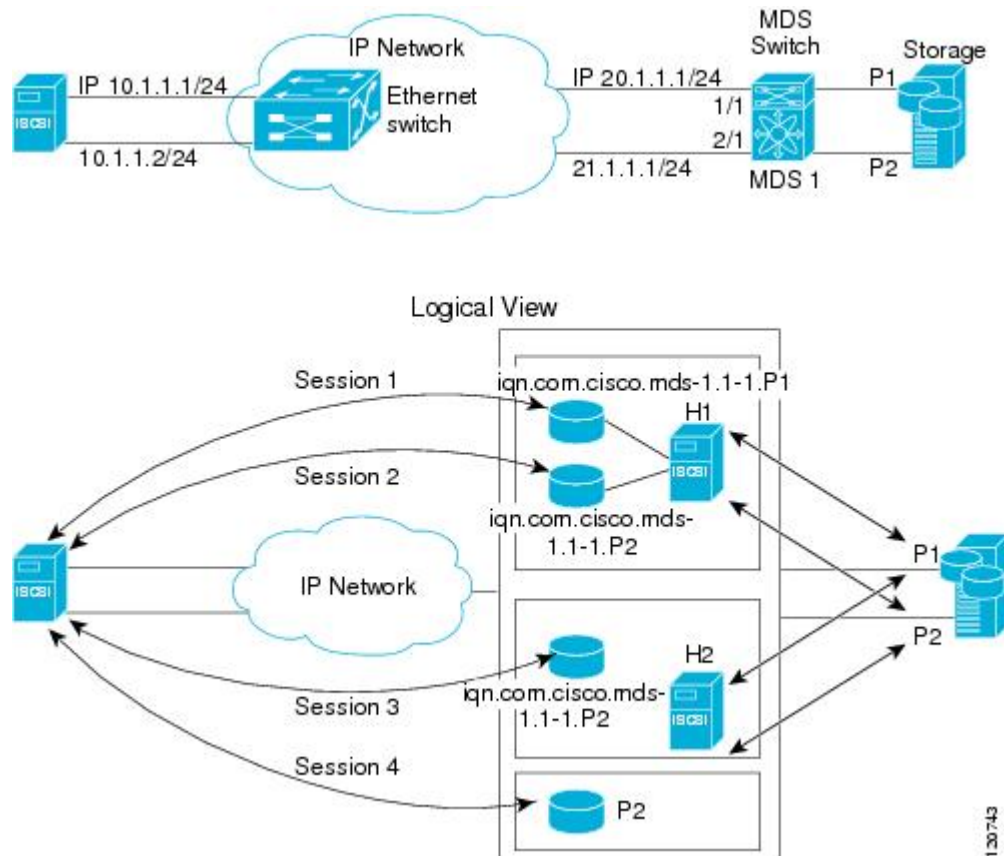
The following high availability features are available for iSCSI configurations:

- iSCSI high availability with host running multi-path software—In this topology, you have recovery from failure of any of the components. The host multi-path software takes care of load balancing or failover across the different paths to access the storage.
- iSCSI high availability with host not having multi-path software—Without multi-path software, the host does not have knowledge of the multiple paths to the same storage.

iSCSI High Availability with Host Running Multi-Path Software

[Figure 131: Host Running Multi-Path Software, on page 826](#) shows the physical and logical topology for an iSCSI HA solution for hosts running multi-path software. In this scenario, the host has four iSCSI sessions. There are two iSCSI sessions from each host NIC to the two IPS ports.

Figure 131: Host Running Multi-Path Software



Each IPS ports is exporting the same two Fibre Channel target ports of the storage but as different iSCSI target names if you use dynamic iSCSI targets). So the two IPS ports are exporting a total of four iSCSI target devices. These four iSCSI targets map the same two ports of the Fibre Channel target.

The iSCSI host uses NIC-1 to connect to IPS port 1 and NIC-2 to connect to IPS port 2. Each IPS port exports two iSCSI targets, so the iSCSI host creates four iSCSI sessions.

If the iSCSI host NIC-1 fails (see [Figure 131: Host Running Multi-Path Software](#), on page 826 for the physical view), then sessions 1 and 2 fail but we still have sessions 3 and 4.

If the IPS port 1 fails, the iSCSI host cannot connect to the IPS port, and sessions 1 and 2 fail. But sessions 3 and 4 are still available.

If the storage port 1 fails, then the IPS ports will terminate sessions 1 and 3 (put iSCSI virtual target iqn.com.cisco.mds-5.1-2.p1 and iqn-com.cisco.mds-5.1-1.p1 in offline state). But sessions 2 and 4 are still available.

In this topology, you have recovery from failure of any of the components. The host multi-path software takes care of load-balancing or failover across the different paths to access the storage.

iSCSI HA with Host Not Having Any Multi-Path Software

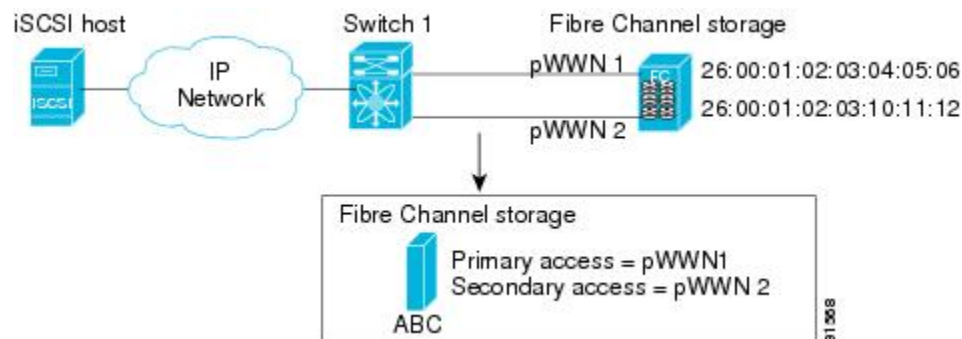
The above topology will not work if the host does not have multi-path software because the host has multiple sessions to the same storage. Without multi-path software the host does not have knowledge of the multiple paths to the same storage.

IP storage has two additional features that provide an HA solution in this scenario.

- IPS ports support the VRRP feature
- IPS has transparent Fibre Channel target failover for iSCSI static virtual targets.

Statically imported iSCSI targets have an additional option to provide a secondary pWWN for the Fibre Channel target. This can be used when the physical Fibre Channel target is configured to have an LU visible across redundant ports. When the active port fails, the secondary port becomes active and the iSCSI session switches to use the new active port (see [Figure 132: Static Target Importing Through Two Fibre Channel Ports, on page 827](#)).

Figure 132: Static Target Importing Through Two Fibre Channel Ports



In [Figure 132: Static Target Importing Through Two Fibre Channel Ports, on page 827](#), you can create an iSCSI virtual target that is mapped to both pWWN1 and pWWN2 to provide redundant access to the Fibre Channel targets.

The failover to a secondary port is done transparently by the IPS port without impacting the iSCSI session from the host. All outstanding I/Os are terminated with a check condition status when the primary port fails. New I/Os received during the failover are not completed and receive a busy status.



Tip If you use LUN mapping, you can define a different secondary Fibre Channel LUN if the LU number is different.

Enable the optional **revert-primary-port** option to direct the IPS port to switch back to the primary port when the primary port is up again. If this option is disabled (default) and the primary port is up again after a switchover, the old sessions will remain with the secondary port and do not switch back to the primary port. However, any new session will use the primary port. This is the only situation when both the primary and secondary ports are used at the same time.


LUN Trespass for Storage Port Failover

In addition to the high availability of statically imported iSCSI targets, the trespass feature is available to enable the move of LUs, on an active port failure, from the active to the passive port of a statically imported iSCSI target.

In physical Fibre Channel targets, which are configured to have LUs visible over two Fibre Channel N ports, when the active port fails, the passive port takes over. Some physical Fibre Channel targets require that the trespass feature be used to move the LUs from the active port to the passive port. A statically imported iSCSI target's secondary pWWN option and an additional option of enabling the trespass feature is available for a physical Fibre Channel target with redundant ports. When the active port fails, the passive port becomes active, and if the trespass feature is enabled, the Cisco MDS switch sends a request to the target to move the

LUs on the new active port. The iSCSI session switches to use the new active port and the moved LUs are accessed over the new active port (see [Figure 133: Virtual Target with an Active Primary Port, on page 828](#)).

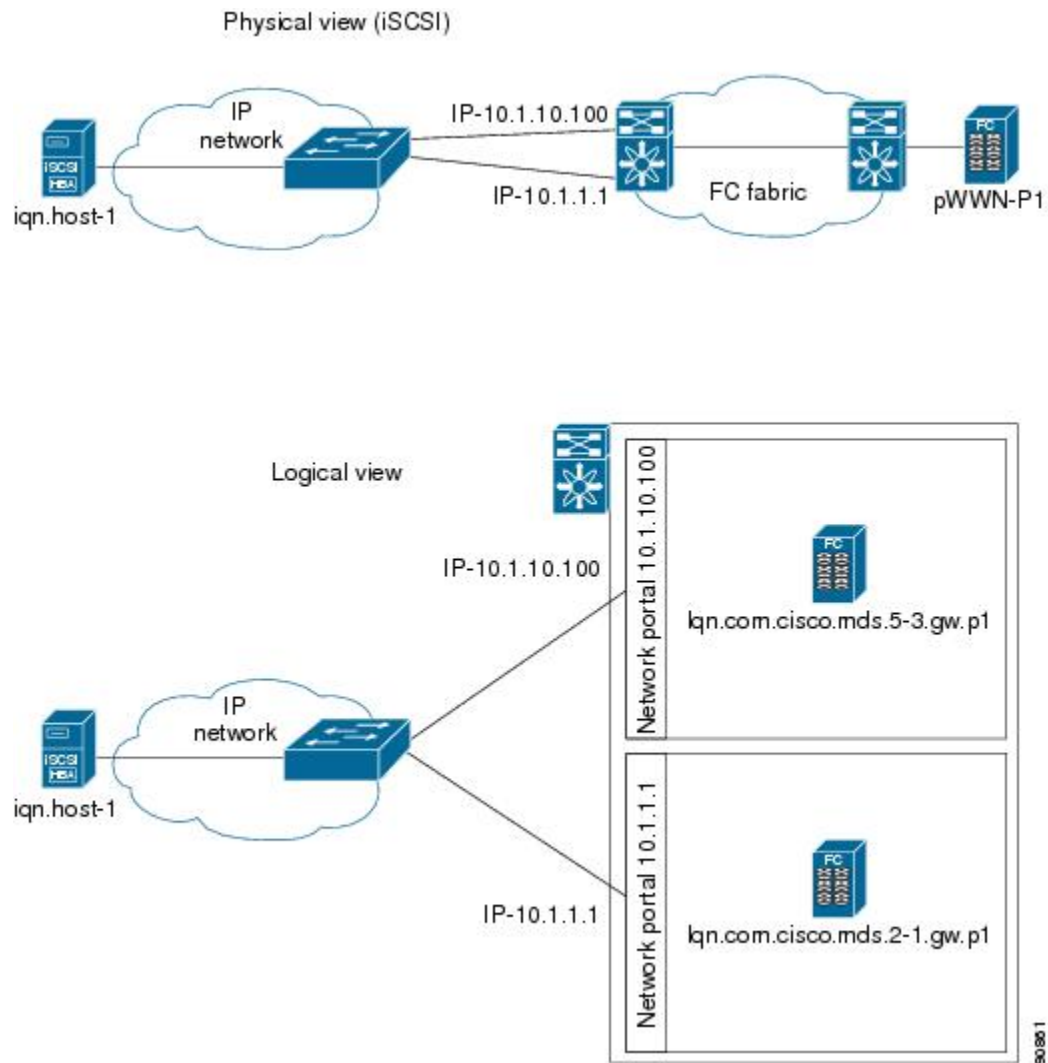
Figure 133: Virtual Target with an Active Primary Port



Multiple IPS Ports Connected to the Same IP Network

[Figure 134: Multiple Gigabit Ethernet Interfaces in the Same IP Network, on page 829](#) provides an example of a configuration with multiple Gigabit Ethernet interfaces in the same IP network.

Figure 134: Multiple Gigabit Ethernet Interfaces in the Same IP Network

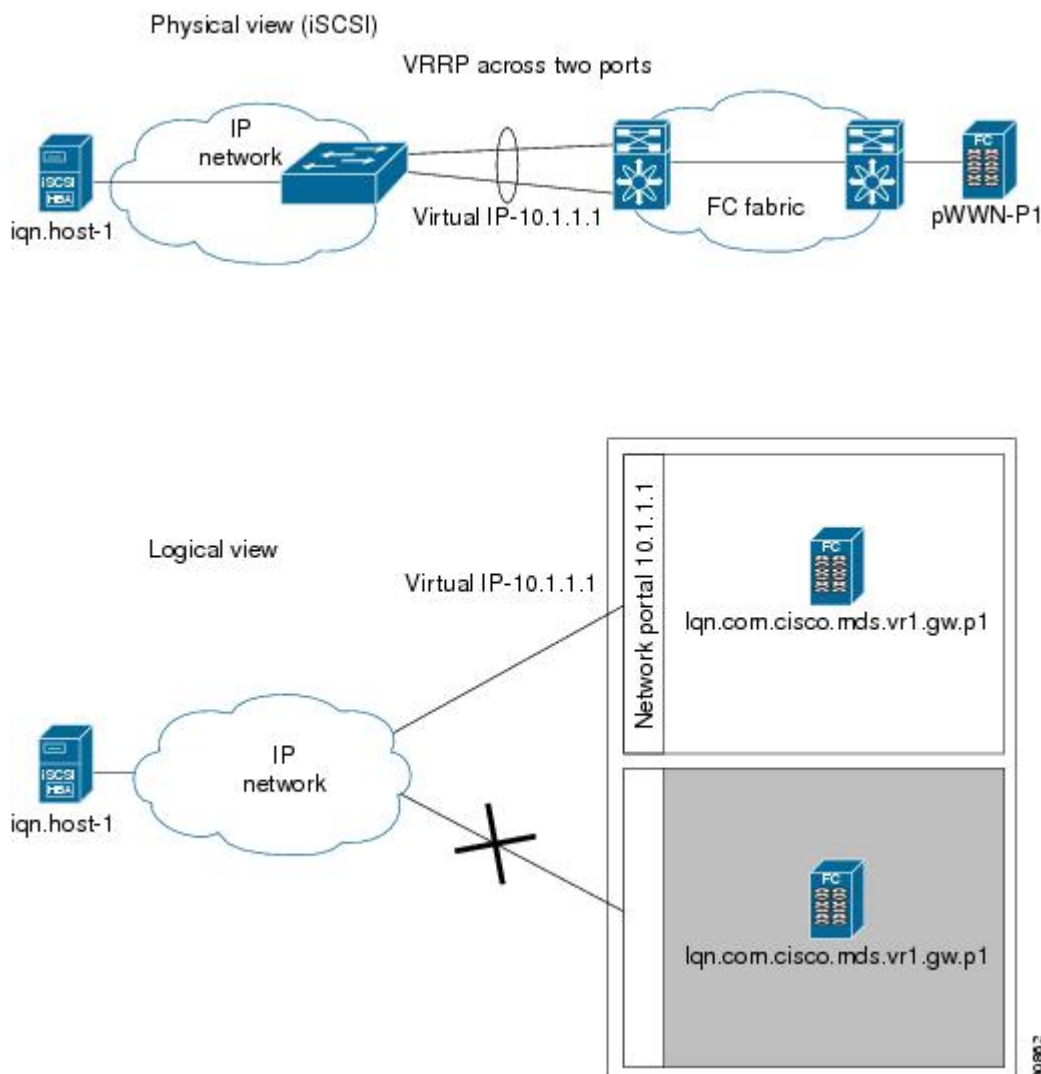


In [Figure 134: Multiple Gigabit Ethernet Interfaces in the Same IP Network, on page 829](#), each iSCSI host discovers two iSCSI targets for every physical Fibre Channel target (with different names). The multi-pathing software on the host provides load-balancing over both paths. If one Gigabit Ethernet interface fails, the host multi-pathing software is not affected because it can use the second path.

VRRP-Based High Availability

[Figure 135: VRRP-Based iSCSI High Availability, on page 830](#) provides an example of a VRRP-based high availability iSCSI configuration.

Figure 135: VRRP-Based iSCSI High Availability




In [Figure 135: VRRP-Based iSCSI High Availability, on page 830](#), each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. When the Gigabit Ethernet interface of the VRRP master fails, the iSCSI session is terminated. The host then reconnects to the target and the session comes up because the second Gigabit Ethernet interface has taken over the virtual IP address as the new master.

Ethernet PortChannel-Based High Availability

All iSCSI data traffic for one iSCSI link is carried on one TCP connection. Consequently, the aggregated bandwidth is 1 Gbps for that iSCSI link.

[Figure 136: Ethernet PortChannel-Based iSCSI High Availability , on page 831](#) provides a sample Ethernet PortChannel-based high availability iSCSI configuration.

Figure 136: Ethernet PortChannel-Based iSCSI High Availability

In [Figure 136: Ethernet PortChannel-Based iSCSI High Availability](#), on page 831, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. The iSCSI session from the iSCSI host to the iSCSI virtual target (on the IPS port) uses one of the two physical interfaces (because an iSCSI session uses one TCP connection). When the Gigabit Ethernet interface fails, the IPS module and the Ethernet switch transparently forwards all the frames on to the second Gigabit Ethernet interface.



Note If an Ethernet PortChannel is configured between the IPS module and an Ethernet switch, the load balancing policy on the Ethernet switch must be based on source/destination IP address only, not port numbers, for load balancing with VRRP to operate correctly.

iSNS

Internet Storage Name Service (iSNS) allows your existing TCP/IP network to function more effectively as a SAN by automating the discovery, management, and configuration of iSCSI devices. To facilitate these functions, the iSNS server and client function as follows:

- The iSNS client registers iSCSI portals and all iSCSI devices accessible through them with an iSNS server.
- The iSNS server provides the following services for the iSNS client:
 - Device registration
 - State change notification
 - Remote domain discovery services

All iSCSI devices (both initiator and target) acting as iSNS clients, can register with an iSNS server. iSCSI initiators can then query the iSNS server for a list of targets. The iSNS server will respond with a list of targets that the querying client can access based on configured access control parameters.

A Cisco MDS 9000 Family switch can act as an iSNS client and register all available iSCSI targets with an external iSNS server. All switches in the Cisco MDS 9000 Family with IPS modules or MPS-14/2 modules installed support iSNS server functionality. This allows external iSNS clients, such as an iSCSI initiator, to register with the switch and discover all available iSCSI targets in the SAN.

About iSNS Client Functionality

Internet Storage Name Service (iSNS) allows your existing TCP/IP network to function more effectively as a SAN by automating the discovery, management, and configuration of iSCSI devices. The iSNS client

registers iSCSI portals and all iSCSI devices accessible through them with an iSNS server. All iSCSI devices (both initiator and target) acting as iSNS clients can register with an iSNS server. When the iSNS client is unable to register or deregister objects with the iSNS server (for example, the client is unable to make a TCP connection to the iSNS server), it retries every minute to reregister all iSNS objects for the affected interfaces with the iSNS server.

The iSNS client functionality on each IPS interface (Gigabit Ethernet interface or subinterface or PortChannel) registers information with an iSNS server.

Once a profile is tagged to an interface, the switch opens a TCP connection to the iSNS server IP address (using the well-known iSNS port number 3205) in the profile and registers network entity and portal objects; a unique entity is associated with each IPS interface. The switch then searches the Fibre Channel name server (FCNS) database and switch configuration to find storage nodes to register with the iSNS server.

Statically mapped virtual targets are registered if the associated Fibre Channel pWWN is present in the FCNS database and no access control configuration prevents it. A dynamically mapped target is registered if dynamic target importing is enabled. See the [Presenting Fibre Channel Targets as iSCSI Targets, on page 808](#) for more details on how iSCSI imports Fibre Channel targets.

A storage node is deregistered from the iSNS server when it becomes unavailable when a configuration changes (such as access control change or dynamic import disabling) or the Fibre Channel storage port goes offline. It is registered again when the node comes back online.

When the iSNS client is unable to register or deregister objects with the iSNS server (for example, the client is unable to make a TCP connection to the iSNS server), it retries every minute to reregister all iSNS objects for the affected interfaces with the iSNS server. The iSNS client uses a registration interval value of 15 minutes. If the client fails to refresh the registration during this interval, the server will deregister the entries.

Untagging a profile also causes the network entity and portal to be deregistered from that interface.


Note

The iSNS client is not supported on a VRRP interface.

About iSNS Server Functionality

When enabled, the iSNS server on the Cisco 9000 Family MDS switch tracks all registered iSCSI devices. As a result, iSNS clients can locate other iSNS clients by querying the iSNS server. The iSNS server also provides the following functionalities:

- Allows iSNS clients to register, deregister, and query other iSNS clients registered with the iSNS server.
- Provides centralized management for enforcing access control to provide or deny access to targets from specific initiators.
- Provides a notification mechanism for registered iSNS clients to receive change notifications on the status change of other iSNS clients.
- Provides a single access control configuration for both Fibre Channel and iSCSI devices.
- Discovers iSCSI targets that do not have direct IP connectivity to the iSCSI initiators.

iSNS Client Registration and Deregistration

You can use the **show isns database** command to display all registered iSNS clients and their associated configuration.

An iSNS client cannot query the iSNS server until it has registered. iSNS client deregistration can occur either explicitly or when the iSNS server detects that it can no longer reach the client (through ESI monitoring).

iSNS client registration and deregistration result in status change notifications (SCNs) being generated to all interested iSNS clients.

Target Discovery

iSCSI initiators discover targets by issuing queries to the iSNS server. The server supports DevGetNext requests to search the list of targets and DevAttrQuery to determine target and portal details, such as the IP address or port number to which to connect.

On receiving a query request from the iSCSI client, the iSNS server queries the Fibre Channel Name Server (FCNS) to obtain a list of Fibre Channel targets that are accessible by the querying initiator. The result of this query depends on zoning configuration currently active and current configuration(s) of the initiator. The iSNS server will subsequently use the iSCSI target configuration(s) (virtual target and dynamic import configuration) to translate the Fibre Channel target to an equivalent iSCSI target. At this stage it also applies any access control configured for the virtual target. A response message with the target details is then sent back to the query initiator.

The iSNS server sends a consolidated response containing all possible targets and portals to the querying initiator. For example, if a Fibre Channel target is exported as different iSCSI targets on different IPS interfaces, the iSNS server will respond with a list of all possible iSCSI targets and portals.

In order to keep the list of targets updated, the iSNS server sends state change notifications (SCN) to the client whenever an iSCSI target becomes reachable or unreachable. The client is then expected to rediscover its list of accessible targets by initiating another iSNS query. Reachability of iSCSI targets changes when any one of the following occurs:

- Target goes up or down.
- Dynamic import of FC target configuration changes.
- Zone set changes.
- Default zone access control changes.
- IPS interface state changes.
- Initiator configuration change makes the target accessible or inaccessible.

About Cloud Discovery

When an iSNS server receives a query request, it responds with a list of available targets and the portals through which the initiator can reach the target. The IP network configuration outside the MDS switch may result in only a subset of Gigabit Ethernet interfaces being reachable from the initiator. To ensure that the set of portals returned to the initiator is reachable, the iSNS server needs to know the set of Gigabit Ethernet interfaces that are reachable from a given initiator.



Note iSNS Cloud Discovery is not supported on the Cisco Fabric Switch for IBM BladeCenter and Cisco Fabric Switch for HP c-Class BladeSystem.

The iSNS cloud discovery feature provides information to the iSNS server on the various interfaces reachable from an initiator by partitioning the interfaces on a switch into disjointed IP clouds. This discovery is achieved by sending messages to all other known IPS ports that are currently up and, depending on the response (or the lack of it), determines if the remote IPS port is in the same IP network or in a different IP network.

Cloud discovery is initiated when the following events occur:

- Manual requests from the CLI initiate cloud discovery from the CLI. This action causes the destruction of existing memberships and makes new ones.
- Auto-discovery of the interface results in an interface being assigned to its correct cloud. All other cloud members are not affected. The membership of each cloud is built incrementally and is initiated by the following events:
 - A Gigabit Ethernet interface comes up. This can be a local or remote Gigabit Ethernet interface.
 - The IP address of a Gigabit Ethernet interface changes.
 - The VRRP configuration on a port changes.

The iSNS server distributes cloud and membership information across all the switches using CFS. Therefore, the cloud membership view is the same on all the switches in the fabric.

**Note**

For CFS distribution to operate correctly for iSNS cloud discovery, all switches in the fabric must be running Cisco SAN-OS Release 3.0(1) or NX-OS 4.1(1b) and later.

Licensing Requirements for iSCSI

The following table shows the licensing requirements for this feature:

License	License Description
Enterprise package (ENTERPRISE_PKG)	It comprises the IP security (IPsec) protocol for iSCSI and FCIP using the MPS-14/2 module or Cisco MDS 9216i Switch.

Guidelines and Limitations

iSLB configuration has the following limits:

- The maximum number of iSLB and iSCSI initiators supported in a fabric is 2000.
- The maximum number of iSLB and iSCSI sessions supported by an IPS port in either transparent or proxy initiator mode is 500.
- The maximum number of iSLB initiators supported in a fabric is 2000.
- The maximum number of iSLB initiators and iSCSI sessions supported by a switch is 5000.
- The maximum number of iSLB sessions per IPS port in either transparent or proxy initiator mode is 500.
- The maximum number of iSLB and iSCSI targets supported in a fabric is 6000.
- The maximum number of switches in a fabric that can have iSLB with CFS distribution enabled is four.
- No more than 200 new iSLB initiators can be added to the pending configuration. Before adding more initiators, you must commit the configuration.
- You cannot disable iSCSI if you have more than 200 iSLB initiators in the running configuration. Reduce the number of iSLB initiators to fewer than 200 before disabling iSCSI.
- iSLB can be used without CFS distribution but if iSLB auto-zone feature is used, traffic is disrupted when any zoneset is activated.
- If IVR and iSLB features are enabled in the same fabric, you should have at least one switch in the fabric where both these features are enabled. Any zoning-related configuration and activation (for normal zones, IVR zones, or iSLB zones) must be performed on this switch. Otherwise, there may be traffic disruption in the fabric.

Default Settings

Table 106: Default iSCSI Parameters , on page 835 lists the default settings for iSCSI parameters.

Table 106: Default iSCSI Parameters

Parameters	Default
Number of TCP connections	One per iSCSI session
minimum-retransmit-time	300 msec
keepalive-timeout	60 seconds
max-retransmissions	4 retransmissions
PMTU discovery	Enabled
pmtu-enable reset-timeout	3600 sec
SACK	Enabled
max-bandwidth	1 Gbps
min-available-bandwidth	70 Mbps
round-trip-time	1 msec
Buffer size	4096 KB
Control TCP and data connection	No packets are transmitted
TCP congestion window monitoring	Enabled
Burst size	50 KB
Jitter	500 microseconds
TCP connection mode	Active mode is enabled
Fibre Channel targets to iSCSI	Not imported
Advertising iSCSI target	Advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces
iSCSI hosts mapping to virtual Fibre Channel hosts	Dynamic mapping
Dynamic iSCSI initiators	Members of the VSAN 1
Identifying initiators	iSCSI node names
Advertising static virtual targets	No initiators are allowed to access a virtual target (unless explicitly configured)
iSCSI login authentication	CHAP or none authentication mechanism

Parameters	Default
revert-primary-port	Disabled
Header and data digest	Enabled automatically when iSCSI initiators send requests. This feature cannot be configured and is not available in store-and-forward mode.
iSNS registration interval	60 sec (not configurable)
iSNS registration interval retries	3
Fabric distribution	Disabled

Table 107: Default iSLB Parameters , on page 836 lists the default settings for iSLB parameters.

Table 107: Default iSLB Parameters

Parameters	Default
Fabric distribution	Disabled
Load balancing metric	1000

Configuring iSCSI

This section describes how to configure iSCSI on the Cisco MDS 9000 Family switches.

This section includes the following sections:

Enabling iSCSI

To use the iSCSI feature, you must explicitly enable iSCSI on the required switches in the fabric. Alternatively, you can enable or disable the iSCSI feature directly on the required modules using Cisco DCNM for SAN or Device Manager. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.



Caution

When you disable this feature, all related configurations are automatically discarded.

Enabling iSCSI on any switch

To enable iSCSI on any switch, follow these steps:

Procedure

Step 1

Choose **FC Services > iSCSI** in the Physical Attributes pane.

You see the iSCSI tables in the Information pane.

The **Control** tab is the default tab. You see the iSCSI enable status for all switches in the fabric that contain IPS ports.

- Step 2** Choose **enable** from the Command column for each switch that you want to enable iSCSI on.
- Step 3** Click the **Apply Changes** icon to save these changes.
-

Enabling iSCSI on a module

To enable iSCSI on a module, follow these steps:

Procedure

- Step 1** Choose **FC Services > iSCSI** in the Physical Attributes pane.
You see the iSCSI tables in the Information pane.
- Step 2** Click the Module Control tab.
You see the Module Control dialog box in the information pane.
- Step 3** Check the Mode Admin check box to enable iSCSI for a specified port on the selected module.
- Step 4** Click the **Apply Changes** icon to save these changes.
-

Enabling iSCSI on a module using Device Manager

To enable iSCSI on a module using Device Manager, follow these steps:

Procedure

- Step 1** Choose **IP > iSCSI**
You see the iSCSI table.
- Step 2** Check the Mode Admin check box to enable iSCSI for the specified port on the selected module.
- Step 3** Click Apply to save these changes.
-

Creating iSCSI Interfaces

Each physical Gigabit Ethernet interface on an IPS module or MPS-14/2 module can be used to translate and route iSCSI requests to Fibre Channel targets and responses in the opposite direction. To enable this capability, the corresponding iSCSI interface must be in an enabled state.

Using the iSCSI Wizard

To use the iSCSI wizard in Cisco DCNM-SAN, follow these steps:

Procedure

- Step 1** Click the **iSCSI > Setup Wizard** icon.

You see the iSCSI Wizard Configure Initiator dialog box.

Step 2 Select an existing iSCSI initiator or add the iSCSI node name or IP address for a new iSCSI initiator.

Step 3 Select the switch for this iSCSI initiator if you are adding a new iSCSI initiator and click **Next**.

You see the iSCSI Wizard Select Targets dialog box.

Step 4 Select the VSAN and targets to associate with this iSCSI initiator and click **Next**.

Note The iSCSI wizard turns on the Dynamic Import FC Targets feature.

You see the iSCSI Wizard Select Zone dialog box.

Step 5 Set the zone name for this new iSCSI zone and check the **ReadOnly** check box if needed.

Step 6 Click **Finish** to create this iSCSI initiator.

If created, the target VSAN is added to the iSCSI host VSAN list.

Note iSCSI wizard automatically turns on the Dynamic FC target import.

Enabling Dynamic Mapping

To enable dynamic mapping of Fibre Channel targets into iSCSI using Device Manager, follow these steps:

Procedure

Step 1 Choose **IP > iSCSI**.

You see the iSCSI configuration.

Step 2 Click the Target tab to display a list of existing iSCSI targets.

Step 3 Check the Dynamically Import FC Targets check box.

Step 4 Click Apply to save this change.

Creating Static Mapping

To create a static iSCSI virtual target for the entire Fibre Channel target port using Device Manager, follow these steps:

Procedure

Step 1 Click **IP > iSCSI**.

You see the iSCSI configuration.

Step 2 Click the Targets tab to display a list of existing iSCSI targets .

Step 3 Click Create to create an iSCSI target.

You see the Create iSCSI Targets dialog box.

- Step 4** Set the iSCSI target node name in the iSCSI Name field, in IQN format.
- Step 5** Set the Port WWN field for the Fibre Channel target port you are mapping.
- Step 6** Click the **Select from List** radio button and set the iSCSI initiator node names or IP addresses that you want this virtual iSCSI target to access, or click the **All** radio button to let the iSCSI target access all iSCSI initiators. Also see the [iSCSI Access Control, on page 814](#).
- Step 7** Click the **Select from List** radio button and check each interface you want to advertise the iSCSI targets on or click the **All** radio button to advertise all interfaces.
- Step 8** Click **Apply** to save this change.

What to do next



Tip An iSCSI target cannot contain more than one Fibre Channel target port. If you have already mapped the whole Fibre Channel target port, you cannot use the LUN mapping option.



Note See the [iSCSI-Based Access Control, on page 815](#) for more information on controlling access to statically mapped targets.

Advertising Static iSCSI Targets

You can limit the Gigabit Ethernet interfaces through which static iSCSI targets are advertised. By default iSCSI targets are advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces.

To configure a specific interface that should advertise the iSCSI virtual target using Device Manager, follow these steps:

Procedure

-
- Step 1** Select **IP > iSCSI**.
You see the iSCSI configuration.
- Step 2** Click the Targets tab to display a list of existing iSCSI targets.
- Step 3** Right-click the iSCSI target that you want to modify and click **Edit Advertised**.
You see the Advertised Interfaces dialog box.
- Step 4** (Optional) Right-click an interface that you want to delete and click **Delete**.
- Step 5** (Optional) Click Create to advertise on more interfaces.
You see the Create Advertised Interfaces dialog box.
-

Specifying the Initiator Identification

You can configure the iSCSI initiator identification mode on each IPS port and all the iSCSI hosts terminating on the IPS port will be identified according to that configuration. The default mode is to identify the initiator by name.

To specify the initiator identification mode, follow these steps:

Procedure

-
- Step 1** Choose **Interfaces > FC Logical** from the Physical Attributes pane.
You see the interfaces configuration in the Information pane.
 - Step 2** Click the **iSCSI** tab.
You see the iSCSI interfaces configuration.
 - Step 3** Right-click the Initiator ID Mode field for the iSCSI interface that you want to modify and select **name** or **ipaddress** from the drop-down menu.
 - Step 4** Click **Apply Changes** to save this change.
-

Configuring the iSCSI Initiator Idle Timeout

iSCSI initiator idle timeout specifies the time for which the virtual Fibre Channel N port is kept idle after the initiator logs out from its last iSCSI session. The default value for this timer is 300 seconds. This is useful to avoid N ports logging in to and logging off of the Fibre Channel SAN as transient failure occurs in the IP network. This helps reduce unnecessary RSCNs being generated in the Fibre Channel SAN.

To configure the initiator idle timeout, follow these steps:

Procedure

-
- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.
You see the iSCSI tables in the Information pane.
 - Step 2** Click the **Globals** tab.
You see the iSCSI global configuration.
 - Step 3** Right-click on the **Initiator Idle Timeout** field that you want to modify and enter the new timeout value.
 - Step 4** Click the **Apply Changes** icon to save these changes.
-

Configuring Static Mapping

To configure static mapping for an iSCSI initiator using Device Manager, follow these steps:

Procedure

-
- Step 1** Select **IP > iSCSI**.
- You see the iSCSI configuration. The Initiators tab is the default.
- Step 2** Click Create to create an iSCSI initiator.
- You see the Create iSCSI Initiators dialog box.
- Step 3** Set the iSCSI node name or IP address and VSAN membership.
- Step 4** In the Node WWN section, check the **Persistent** check box.
- Step 5** Check the **System Assigned** check box if you want the switch to assign the nWWN or leave this unchecked and set the Static WWN field.
- Step 6** In the Port WWN section, check the **Persistent** check box if you want to statically map pWWNs to the iSCSI initiator.
- Step 7** If persistent, check the **System Assigned** check box and set the number of pWWNs to reserve for this iSCSI initiator if you want the switch to assign pWWNs. Alternately, you can leave this unchecked and set one or more pWWNs for this iSCSI initiator.
- Step 8** (Optional) Set the AuthUser field if authentication is enabled. Also see the [iSCSI Session Authentication, on page 816](#).
- Step 9** Click **Create** to create this iSCSI initiator.

Note If the system-assign option is used to configure WWNs for an iSCSI initiator, when the configuration is saved to an ASCII file the system-assigned WWNs are also saved. Subsequently if you perform a write erase, you must manually delete the WWN configuration from the ASCII file. Failing to do so can cause duplicate WWN assignments if the ASCII configuration file is reapplied on the switch.

Making the Dynamic iSCSI Initiator WWN Mapping Static

After a dynamic iSCSI initiator has already logged in, you may decide to permanently keep the automatically assigned nWWN/pWWN mapping so this initiator uses the same mapping the next time it logs in.

You can convert a dynamic iSCSI initiator to static iSCSI initiator and make its WWNs persistent.



Note You cannot convert a dynamic iSCSI initiator to a static iSLB initiator or a dynamic iSLB initiator to a static iSCSI initiator.



Note Making the dynamic pWWNs static after the initiator is created is supported only through the CLI, not through Device Manager or Cisco DCNM- SAN. In Cisco DCNM-SAN or Device Manager, you must delete and then recreate this initiator to have the pWWNs static.

Checking for WWN Conflicts

WWNs assigned to static iSCSI initiators by the system can be inadvertently returned to the system when an upgrade fails or you downgrade the system software (manually booting up an older Cisco MDS SAN-OS release without using the **install all** command). In these instances, the system can later assign those WWNs to other iSCSI initiators (dynamic or static) and cause conflicts.

You can address this problem by checking for and removing any configured WWNs that belong to the system whenever such scenarios occur.

To permanently keep the automatically assigned nWWN mapping, follow these steps:

Procedure

-
- Step 1** Choose **End Devices** > **iSCSI** in the Physical Attributes pane.
You see the iSCSI tables in the Information pane.
 - Step 2** Click the **Initiators** tab.
You see the iSCSI initiators configured.
 - Step 3** Check the **PersistentNodeWWN** check box for the iSCSI initiators that you want to make static.
 - Step 4** Click the **ApplyChanges** icon to save these changes.
-

Configuring the Proxy Initiator

To configure the proxy initiator, follow these steps:

Procedure

-
- Step 1** Expand **Switches**, expand FC **Interfaces**, and then select **Logical** in the Physical Attributes pane.
You see the Interface tables in the Information pane.
 - Step 2** In Device Manager, select Interface > **Ethernet and iSCSI**.
You see the Ethernet Interfaces and iSCSI dialog box.
 - Step 3** Click the **iSCSI** tab in either FM or DM.
You see the iSCSI interface configuration table.
 - Step 4** Check the **Proxy Mode Enable** check box.
 - Step 5** Click the **Apply Changes** icon in Cisco DCNM-SAN or click Apply in Device Manager to save these changes.
-

What to do next



Note When an interface is in proxy initiator mode, you can only configure Fibre Channel access control (zoning) based on the iSCSI interface's proxy N port attributes—the WWN pairs or the FC ID. You cannot configure zoning using iSCSI attributes such as IP address or IQN of the iSCSI initiator. To enforce initiator-based access control, use iSCSI based access control (see the [iSCSI Access Control, on page 814](#)).

Configuring VSAN Membership for iSCSI Hosts

Individual iSCSI hosts can be configured to be in a specific VSAN. The specified VSAN overrides the iSCSI interface VSAN membership.

To assign VSAN membership for iSCSI hosts, follow these steps:

Procedure

- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.
You see the iSCSI tables in the Information pane.
- Step 2** Click the **Initiators** tab.
You see the iSCSI initiators configured.
- Step 3** Fill in the VSAN Membership field to assign a VSAN to the iSCSI hosts.
- Step 4** Click the **Apply Changes** icon to save these changes.

Note When an initiator is configured in any other VSAN (other than VSAN 1), for example VSAN 2, the initiator is automatically removed from VSAN 1. If you also want it to be present in VSAN 1, you must explicitly configure the initiator in VSAN 1.

Configuring Default Port VSAN for iSCSI Interfaces

VSAN membership can be configured for an iSCSI interface, called the *port VSAN*. All the iSCSI devices that connect to this interface automatically become members of this VSAN, if it is not explicitly configured in a VSAN. In other words, the port VSAN of an iSCSI interface is the default VSAN for all dynamic iSCSI initiators. The default port VSAN of an iSCSI interface is VSAN 1.



Caution Changing the VSAN membership of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the [Changing iSCSI Interface Parameters and the Impact on Load Balancing, on page 823](#).

To change the default port VSAN for an iSCSI interface using Device Manager, follow these steps:

Procedure

- Step 1** Choose **Interface > Ethernet and iSCSI**.
You see the Ethernet Interfaces and iSCSI dialog box.
- Step 2** Click the iSCSI tab.
You see the iSCSI interface configuration table.
- Step 3** Double-click the PortVSAN column and modify the default port VSAN.
- Step 4** Click **Apply** to save these changes.
-

Adding iSCSI Initiator to the Zone Database

To add an iSCSI initiator to the zone database, follow these steps:

Procedure

- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Edit Local Zone Database dialog box.
- Step 2** Select the VSAN you want to add the iSCSI host initiator to and click **OK**.
You see the available zones and zone sets for that VSAN.
- Step 3** From the list of available devices with iSCSI host initiators, drag the initiators to add into the zone.
- Step 4** Click **Distribute** to distribute the change.
-

Configuring Access Control in iSCSI

To configure access control in iSCSI using Device Manager, follow these steps:

Procedure

- Step 1** Select **IP > iSCSI**.
You see the iSCSI configuration.
- Step 2** Click the **Targets** tab.
You see the iSCSI virtual targets.
- Step 3** Uncheck the **Initiators Access All** check box if checked.
- Step 4** Click **Edit Access**.
You see the Initiators Access dialog box.

- Step 5** Click **Create** to add more initiators to the Initiator Access list.
You see the Create Initiators Access dialog box.
- Step 6** Add the name or IP address for the initiator that you want to permit for this virtual target.
- Step 7** Click **Create** to add this initiator to the Initiator Access List.
-

Configuring AAA Authentication for an iSCSI User

To configure AAA authentication for an iSCSI user, follow these steps:

Procedure

- Step 1** Choose **Switches > Security > AAA** in the Physical Attributes pane.
You see the AAA configuration in the Information pane.
- Step 2** Click the **Applications** tab.
You see the AAA configuration per application.
- Step 3** Right-click the **ServerGroup Id List** field for the iSCSI application and enter the server group that you want iSCSI to use.
- Note** You should use an existing server group or create a new server group before configuring it for iSCSI session authentication.
- Step 4** Click the **Apply Changes** icon to save these changes.
-

Configuring Authentication Mechanism

You can configure iSCSI CHAP or None authentication at both the global level and at each interface level.

The authentication for a Gigabit Ethernet interface or subinterface overrides the authentication method configured at the global level.

If CHAP authentication is used, issue the **iscsi authentication chap** command at either the global level or at a per-interface level. If authentication should not be used at all, issue the **iscsi authentication none** command.

Configuring AAA authentication for an iSCSI user

To configure AAA authentication for an iSCSI user, follow these steps:

Procedure

- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.
You see the iSCSI tables in the Information pane.
- Step 2** Click the **Globals** tab.

You see the iSCSI authentication configuration table.

Step 3 Select **chap** or **none** from the authMethod column.

Step 4 Click the **Apply > Changes** icon in Cisco DCNM-SAN to save these changes.

Configuring the authentication mechanism for iSCSI sessions

To configure the authentication mechanism for iSCSI sessions to a particular interface, follow these steps:

Procedure

Step 1 Choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.

You see the Gigabit Ethernet configuration in the Information pane.

Step 2 Click the iSNS tab.

You see the iSCSI and iSNS configuration.

Step 3 Right-click on the **IscsiAuthMethod** field and select none or chap.

Step 4 Click the **Apply Changes** icon to save these changes.

Configuring Local Authentication

See the Security Configuration Guide, Cisco DCNM for SAN to create the local password database. To create users in the local password database for the iSCSI initiator, the iSCSI keyword is mandatory.

To configure iSCSI users for local authentication using Device Manager, follow these steps:

Procedure

Step 1 Choose **Security > iSCSI**.

You see the iSCSI Security dialog box.

Step 2 Complete the iSCSI User, Password, and Password Confirmation fields.

Step 3 Click **Create** to save this new user.

Restricting iSCSI Initiator Authentication

By default, the iSCSI initiator can use any user name in the RADIUS server or in the local database in authenticating itself to the IPS module or MPS-14/2 module (the CHAP user name is independent of the iSCSI initiator name). The IPS module or MPS-14/2 module allows the initiator to log in as long as it provides a correct response to the CHAP challenge sent by the switch. This can be a problem if one CHAP user name and password has been compromised.

To restrict an initiator to use a specific user name for CHAP authentication, follow these steps:

Procedure

- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.
You see the iSCSI tables in the Information pane.
- Step 2** Right-click the AuthUser field and enter the user name to which you want to restrict the iSCSI initiator.
- Step 3** Click the **Apply Changes** icon to save these changes.
-

Configuring Mutual CHAP Authentication

The IPS module or MPS-14/2 module supports a mechanism by which the iSCSI initiator can authenticate the Cisco MDS switch's iSCSI target during the iSCSI login phase. This authentication is available in addition to the IPS module or MPS-14/2 module authentication of the iSCSI initiator.

In addition to the IPS module or MPS-14/2 module authentication of the iSCSI initiator, the IPS module or MPS-14/2 module also supports a mechanism for the iSCSI initiator to authenticate the Cisco MDS switch's iSCSI target during the iSCSI login phase. This authentication requires the user to configure a user name and password for the switch to present to the iSCSI initiator. The provided password is used to calculate a CHAP response to a CHAP challenge sent to the IPS port by the initiator.

Configuring a global iSCSI target user name and password

To configure a global iSCSI target user name and password to be used by the switch to authenticate itself to an initiator, follow these steps:

Procedure

- Step 1** Choose **FC Interfaces > Logical > iSCSI** in the Physical Attributes pane.
You see the iSCSI tables in the Information pane.
- Step 2** Select the **Globals** tab.
You see the global iSCSI configuration.
- Step 3** Fill in the Target UserName and Target Password fields.
- Step 4** Click the **Apply Changes** icon to save these changes.
-

Configuring a per-initiator iSCSI target user name and password

To configure a per-initiator iSCSI target's user name and password used by the switch to authenticate itself to an initiator using Device Manager, follow these steps:

Procedure

- Step 1** Choose **IP > iSCSI**.
You see the iSCSI configuration.

- Step 2** Complete the **Target UserName** and **Target Password** fields for the initiator that you want to configure.
- Step 3** Click **Create** to add this initiator to the Initiator Access List.

Configuring an iSCSI RADIUS Server

To configure an iSCSI RADIUS server, follow these steps:

Procedure

- Step 1** Configure the RADIUS server to allow access from the Cisco MDS switch's management Ethernet IP address.
- Step 2** Configure the shared secret for the RADIUS server to authenticate the Cisco MDS switch.
- Step 3** Configure the iSCSI users and passwords on the RADIUS server.

Setting QoS Values

To set the QoS values, follow these steps:

Procedure

- Step 1** Expand **Switches**, expand **FC Interfaces**, and then select **Logical** in the Physical Attributes pane.
You see the Interface tables in the Information pane.
- Step 2** In Device Manager, choose **Interface > Ethernet and iSCSI**.
You see the Ethernet Interfaces and iSCSI dialog box.
- Step 3** Click the **iSCSI TCP** tab in either Cisco DCNM-SAN or Device Manager.
You see the iSCSI TCP configuration table.
- Step 4** Set the QoS field from 1 to 6.
- Step 5** Click the **Apply Changes** icon in Cisco DCNM-SAN or click **Apply** in Device Manager to save these changes.

Configuring iSLB



Note For iSLB, all switches in the fabric must be running Cisco MDS SAN-OS Release 2.1(1a) or later.

This section covers the following topics:

Configuring iSLB Using Device Manager

Perform the following actions prior to configuring iSLB:

- Enable iSCSI (see the [Enabling iSCSI, on page 836](#) for more information).
- Configure the Gigabit Ethernet interfaces (see the “Configuring Gigabit Ethernet Interface” section on [page 43-5](#)).
- Configure the VRRP groups (see the [Configuring Load Balancing Using VRRP, on page 853](#)).
- Configure and activate a zone set.
- Enable CFS distribution for iSLB (see the [Enabling iSLB Configuration Distribution, on page 853](#)).

To configure iSLB using Device Manager, follow these steps:

Procedure

- Step 1** Choose **IP > iSCSI iSLB**.
You see the iSCSI iSLB dialog box.
- Step 2** Click **Create** to create a new iSCSI iSLB initiator.
You see the Create iSCSI iSLB Initiators dialog box.
- Step 3** Set the Name or IP Address field to the iSLB name or IP address.
- Step 4** Set the VSAN Membership field to the VSAN that you want the iSLB initiator in.
Also see the [Assigning VSAN Membership for iSLB Initiators, on page 850](#).
- Step 5** Check the Persistent check box to convert a dynamic nWWN to static for the iSLB initiator.
Also see the [Making the Dynamic iSCSI Initiator WWN Mapping Static, on page 841](#).
- Step 6** (Optional) Check the **SystemAssigned** check box to have the switch assign the nWWN.
- Step 7** (Optional) Set the Static WWN field to manually assign the static nWWN. You must ensure uniqueness for this nWWN.
- Step 8** (Optional) Check the Port WWN Mapping **Persistent** check box to convert dynamic pWWNs to static for the iSLB initiator.
See the [Making the Dynamic iSCSI Initiator WWN Mapping Static, on page 841](#).
- Step 9** (Optional) Check the **SystemAssigned** check box and set the number of pWWNs you want to have the switch assign the PWWN.
- Step 10** (Optional) Set the Static WWN(s) field to manually assign the static pWWNs.
You must ensure uniqueness for these pWWN.
- Step 11** (Optional) Set the AuthUser field to the username that you want to restrict the iSLB initiator to for iSLB authentication.
Also see the [Restricting iSLB Initiator Authentication , on page 852](#).
- Step 12** Fill in the Username and Password fields to configure iSLB initiator target CHAP authentication.
Also see the [iSLB Session Authentication, on page 821](#).
- Step 13** In the Initiator Specific Target section, set the pWWN to configure an iSLB initiator target.

- Step 14** (Optional) Set the Name field to a globally unique identifier (IQN).
- Step 15** (Optional) Check the **NoAutoZoneCreation** check box to disable auto-zoning.
- Step 16** (Optional) Check the **TresspassMode** check box.
- Also see the [LUN Trespass for Storage Port Failover, on page 827](#).
- Step 17** (Optional) Check the **RevertToPrimary** check box to revert back to the primary port after an HA failover when the primary port comes back up.
- Step 18** Set the PrimaryVsan to the VSAN for the iSLB initiator target.
- Step 19** Click **Create** to create this iSLB initiator.
- Step 20** If CFS is enabled, select **commit** from the CFS drop-down menu.

Configuring iSLB Initiator Names or IP Addresses

You must specify the iSLB initiator name or IP address before configuring it.



Note Specifying the iSLB initiator name or IP address is the same as for an iSCSI initiator. See the [Static Mapping, on page 813](#).

Making the Dynamic iSLB Initiator WWN Mapping Static

After a dynamic iSLB initiator has logged in, you may decide to permanently keep the automatically assigned nWWN/pWWN mapping to allow this initiator to use the same mapping the next time it logs in (see the [Dynamic Mapping, on page 808](#)).

You can convert a dynamic iSLB initiator to a static iSLB initiator and make its WWNs persistent.

- You cannot convert a dynamic iSCSI initiator to a static iSLB initiator or a dynamic iSLB initiator to a static iSCSI initiator (see the “Dynamic Mapping” section on page 4-20) .
- Making the dynamic mapping for iSLB initiators static is the same as for iSCSI. See the [Making the Dynamic iSCSI Initiator WWN Mapping Static, on page 841](#).
- Only statically mapped iSLB initiator configuration is distributed throughout the fabric using CFS. Dynamically and statically configured iSCSI initiator configurations are not distributed.

See the [Configuring iSLB Using Device Manager, on page 848](#).

Assigning VSAN Membership for iSLB Initiators

Individual iSLB hosts can be configured to be in a specific VSAN (similar to the DPVM feature for Fibre Channel). The specified VSAN overrides the iSCSI interface VSAN membership.

For more information, see the Fabric Configuration Guide, Cisco DCNM for SAN.



Note Specifying the iSLB initiator VSAN is the same as for an iSCSI initiator. See the [VSAN Membership for iSCSI, on page 814](#).



Note When an iSLB initiator is configured in any other VSAN (other than VSAN 1, the default VSAN), for example VSAN 2, the initiator is automatically removed from VSAN 1. If you also want it to be present in VSAN 1, you must explicitly configure the initiator in VSAN 1.

See the [Configuring iSLB Using Device Manager, on page 848](#).

Configuring Metric for Load Balancing

You can assign a load metric to each initiator for weighted load balancing. The load calculated is based on the number of initiators on a given iSCSI interface. This feature accommodates initiators with different bandwidth requirements. For example, you could assign a higher load metric to a database server than to a web server. Weighted load balancing also accommodates initiators with different link speeds.

Also, you can configure initiator targets using the device alias or the pWWN. If you configure an IQN for an initiator target, then that name is used to identify the initiator target. Otherwise, a unique IQN is generated for the initiator target.

For more information on load balancing, see the [About Load Balancing Using VRRP, on page 821](#).

Choose **IP > iSCSI iSLB** in Device Manager and set the LoadMetric field to change the load balancing metric for an iSLB initiator.

See the [Configuring iSLB Using Device Manager, on page 848](#).

Configuring iSLB Initiator Targets

You can configure initiator targets using the device alias or the pWWN. You can also optionally specify one or more of the following optional parameters:

- Secondary pWWN
- Secondary device alias
- LUN mapping
- IQN
- VSAN identifier



Note The VSAN identifier is optional if the target is online. If the target is not online, the VSAN identifier is required.

In addition, you can disable auto-zoning.

If you configure an IQN for an initiator target, then that name is used to identify the initiator target. Otherwise, a unique IQN is generated for the initiator target.

To configure additional iSLB initiator targets using Device Manager, follow these steps:

Procedure

Step 1

Choose **IP > iSCSI iSLB**.

You see the iSCSI iSLB dialog box.

- Step 2** Click on the initiator you want to add targets to and click **Edit Initiator Specific Targets**.
You see the Initiator Specific Target dialog box.
- Step 3** Click **Create** to create a new initiator target.
You see the Create Initiator Specific Target dialog box.
- Step 4** Fill in the pWWN field with the initiator target pWWN.
- Step 5** (Optional) Set the Name field to a globally unique identifier (IQN).
- Step 6** (Optional) Check the **NoAutoZoneCreation** check box to disable auto-zoning.
- Step 7** (Optional) Check the **TresspassMode** check box. See the [LUN Trespass for Storage Port Failover, on page 827](#).
- Step 8** (Optional) Check the **RevertToPrimary** check box to revert back to the primary port after an HA failover when the primary port comes back up.
- Step 9** Set the PrimaryVsan to the VSAN for the iSLB initiator target.
- Step 10** Click **Create** to create this iSLB initiator target.
- Step 11** If CFS is enabled, select **commit** from the CFS drop-down menu.

Configuring and Activating Zones for iSLB Initiators and Initiator Targets

You can configure a zone name where the iSLB initiators and initiator targets are added. If you do not specify a zone name, the IPS manager creates one dynamically.

iSLB zone sets have the following restrictions:

- Auto-zoning of the initiator with the initiator targets is enabled by default.
- A zone set must be active in a VSAN for auto-zones to be created in that VSAN.
- iSLB zone set activation might fail if another zone set activation is in process or if the zoning database is locked. Retry the iSLB zone set activation if a failure occurs. To avoid this problem, only perform only one zoning related operation (normal zones, IVR zones, or iSLB zones) at a time.
- Auto-zones are created when the zone set is activated and there has been at least one change in the zoneset. The activation has no effect if only the auto-zones have changed.



Caution

If IVR and iSLB are enabled in the same fabric, at least one switch in the fabric must have both features enabled. Any zoning related configuration or activation operation (for normal zones, IVR zones, or iSLB zones) must be performed on this switch. Otherwise, traffic might be disrupted in the fabric.

Choose **IP > iSCSI iSLB** in Device Manager and set the autoZoneName field to change the auto zone name for an iSLB initiator.

See the [Configuring iSLB Using Device Manager, on page 848](#).

Restricting iSLB Initiator Authentication

By default, the iSLB initiator can use any user name in the RADIUS or local AAA database in authenticating itself to the IPS module or MPS-14/2 module (the CHAP user name is independent of the iSLB initiator name). The IPS module or MPS-14/2 module allows the initiator to log in as long as it provides a correct response

to the CHAP challenge sent by the switch. This can be a problem if one CHAP user name and password have been compromised.

Choose **IP > iSCSI iSLB** in Device Manager and set the AuthName field to restrict an initiator to use a specific user name for CHAP authentication.

See the [Configuring iSLB Using Device Manager, on page 848](#).

Mutual CHAP Authentication

In addition to the IPS module and MPS-14/2 module authentication of the iSLB initiator, the IPS module and MPS-14/2 module also support a mechanism for the iSLB initiator to authenticate the Cisco MDS switch's initiator target during the iSCSI login phase. This authentication requires the user to configure a user name and password for the switch to present to the iSLB initiator. The provided password is used to calculate a CHAP response to a CHAP challenge sent to the IPS port by the initiator.

Choose **IP > iSCSI iSLB** in Device Manager and set the Target Username and Target Password fields to configure a per-initiator user name and password used by the switch to authenticate itself to an initiator.

See the [Configuring iSLB Using Device Manager, on page 848](#).

Configuring Load Balancing Using VRRP

You must first configure VRRP on the Gigabit Ethernet interfaces on the switch that connect to the IP network before configuring VRRP for iSLB.

To configure VRRP load balancing using Device Manager, follow these steps:

Procedure

-
- | | |
|---------------|-----------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Choose IP > iSCSI iSLB .
You see the iSCSI iSLB dialog box. |
| Step 2 | Click the VRRP tab. |
| Step 3 | Click Create to configure VRRP load balancing for iSLB initiators.
You see the Create iSCSI iSLB VRRP dialog box. |
| Step 4 | Set the Vrid to the VRRP group number. |
| Step 5 | Select either ipv4 or ipv6 and check the LoadBalance check box. |
| Step 6 | Click Create to enable load balancing. |
| Step 7 | If CFS is enabled, select commit from the CFS drop-down menu. |
-

Distributing the iSLB Configuration Using CFS

This section contains the following:

Enabling iSLB Configuration Distribution

To enable CFS distribution of the iSLB configuration using Device Manager, follow these steps:

Procedure

-
- Step 1** Choose **Admin > CFS**.
You see the CFS dialog box.
- Step 2** Set the Command field to **enable** for the iSLB feature.
- Step 3** Click **Apply** to save this change.
-

Committing Changes to the Fabric

To apply the pending iSLB configuration changes to the active configuration and to other MDS switches in the fabric, you must commit the changes. The pending configuration changes are distributed and, on a successful commit, the configuration changes are applied to the active configuration in the MDS switches throughout the fabric, the automatic zones are activated, and the fabric lock is released.

To commit iSLB configuration changes to other MDS switches in the fabric, activate iSLB automatic zones, and release the fabric lock using Device Manager, follow these steps:

Procedure

-
- Step 1** Choose **Admin > CFS**.
You see the CFS Configuration dialog box.
- Step 2** Set the Command field to **commit** for the iSLB feature.
- Step 3** Click **Apply** to save this change.
-

Discarding Pending Changes

At any time, you can discard the pending changes to the iSLB configuration and release the fabric lock. This action has no affect on the active configuration on any switch in the fabric.

To discard the pending iSLB configuration changes and release the fabric lock using Device Manager, follow these steps:

Procedure

-
- Step 1** Choose **Admin > CFS**.
You see the CFS Configuration dialog box.
- Step 2** Set the Command field to **abort** for the iSLB feature.
- Step 3** Click **Apply** to save this change.
-

Clearing a Fabric Lock

If you have performed an iSLB configuration task and have not released the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your pending changes are discarded and the fabric lock is released.

The pending changes are only available in the volatile directory and are discarded if the switch is restarted.

To release a fabric lock using Device Manager, follow these steps:

Procedure

- Step 1** Choose **Admin > CFS**.
You see the CFS Configuration dialog box.
- Step 2** Set the Command field to **clear** for the iSLB feature.
- Step 3** Click **Apply** to save this change.
-

Creating a Static iSCSI Virtual Target

To create a static iSCSI virtual target for the entire Fibre Channel target port using Device Manager, follow these steps:

Procedure

- Step 1** Click **IP > iSCSI**.
You see the iSCSI configuration.
- Step 2** Click the **Targets** tab to display a list of existing iSCSI targets shown.
- Step 3** Click **Create** to create an iSCSI target.
You see the Create iSCSI Targets dialog box.
- Step 4** Set the iSCSI target node name in the iSCSI Name field, in IQN format.
- Step 5** Set the Port WWN field for the Fibre Channel target port you are mapping.
- Step 6** Click the **Select from List** radio button and set the iSCSI initiator node names or IP addresses that you want this virtual iSCSI target to access, or click the **All** radio button to let the iSCSI target access all iSCSI initiators. See the [iSCSI Access Control, on page 814](#).
- Step 7** Click the **Select from List** radio button and check each interface you want to advertise the iSCSI targets on or choose the **All** radio button to advertise all interfaces.
- Step 8** Click **Apply** to save this change.
-

Enabling the Trespass Feature for a Static iSCSI

In Device Manager, choose **IP > iSCSI**, select the **Targets** tab, and check the **Trespass Mode** check box to enable the trespass feature for a static iSCSI virtual target.

Configuring iSCSI Authentication

This section provides configuration information on iSCSI authentication. It includes the following authentication procedures:



Note

This section does not specify the steps to enter or exit EXEC mode, configuration mode, or any submode. Be sure to verify the prompt before entering any command.



Caution

Changing the authentication of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the [Changing iSCSI Interface Parameters and the Impact on Load Balancing](#), on page 823.

Configuring No Authentication

To configure a network with no authentication, set the iSCSI authentication method to **none**.

In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane. Select the **Globals** tab and set the AuthMethod drop-down menu to **none** and click **Apply Changes**.

Configuring CHAP with Local Password Database

To configure authentication using the CHAP option with the local password database, follow these steps:

Procedure

- Step 1** Set the AAA authentication to use the local password database for the iSCSI protocol:
 - a) In Cisco DCNM-SAN, choose **Switches > Security > AAA** in the Physical Attributes pane.
 - b) Click the **Applications** tab in the Information pane.
 - c) Check the **Local** check box for the iSCSI row and click **Apply Changes**.
- Step 2** Set the iSCSI authentication method to require CHAP for all iSCSI clients:
 - a) In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane.
 - b) Click the **Globals** tab in the Information pane.
 - c) Set the AuthMethod drop-down menu to **chap** and click **Apply Changes**.
- Step 3** Configure the user names and passwords for iSCSI users:
 - a) In Device Manager, choose **Security > iSCSI**.
 - b) Set the Username, Password and Confirm Password fields.
 - c) Click **Create** to save these changes.
- Step 4** Verify the global iSCSI authentication setup:
 - a) In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane.
 - b) Click the **Globals** tab in the Information pane.

Configuring CHAP with External RADIUS Server

Procedure

-
- Step 1** Configure the password for the Cisco MDS switch as RADIUS client to the RADIUS server:
- In Cisco DCNM-SAN, choose **Switches > Security > AAA > RADIUS** in the Physical Attributes pane.
 - Click the **Default** tab in the Information pane.
 - Set the AuthKey field to the default password and click the **Apply Changes** icon.
- Step 2** Configure the RADIUS server IP address:
- In Cisco DCNM-SAN, choose **Switches > Security > AAA > RADIUS** in the Physical Attributes pane.
 - Click the **Server** tab in the Information pane and click **Create Row**.
 - Set the Index field to a unique number.
 - Set the IP Type radio button to **ipv4** or **ipv6**.
 - Set the Name or IP Address field to the IP address of the RADIUS server and click **Create**.
- Step 3** Create a RADIUS server group and add the RADIUS server to the group:
- In Cisco DCNM-SAN, choose **Switches > Security > AAA** in the Physical Attributes pane.
 - Select the **Server Groups** tab in the Information pane and click **Create Row**.
 - Set the Index field to a unique number.
 - Set the Protocol radio button to **radius**.
 - Set the Name field to the server group name.
 - Set the ServerIDList to the index value of the RADIUS server and click **Create**.
- Step 4** Set up the authentication verification for the iSCSI protocol to go to the RADIUS server.
- In Cisco DCNM-SAN, choose **Switches > Security > AAA** in the Physical Attributes pane.
 - Click the **Applications** tab in the Information pane.
 - Right-click on the iSCSI row in the Type, SubType, Function column.
 - Set the ServerGroup IDList to the index value of the Server Group and click **Create**.
- Step 5** Set up the iSCSI authentication method to require CHAP for all iSCSI clients.
- In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane.
 - Select **chap** from the AuthMethod drop-down menu.
 - Click the **Apply Changes** icon.
- Step 6** In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane.
- Step 7** Click the **Globals** tab in the Information pane to verify that the global iSCSI authentication setup is for CHAP.
- Step 8** In Cisco DCNM-SAN, choose **Switches > Security > AAA** in the Physical Attributes pane.
- Step 9** Click the **Applications** tab in the Information pane to verify the AAA authentication information for iSCSI.
-

Configuring an iSCSI RADIUS server

To configure an iSCSI RADIUS server, follow these steps:

Procedure

- Step 1** Configure the RADIUS server to allow access from the Cisco MDS switch's management Ethernet IP address.
 - Step 2** Configure the shared secret for the RADIUS server to authenticate the Cisco MDS switch.
 - Step 3** Configure the iSCSI users and passwords on the RADIUS server.
-

Creating an iSNS Client Profile

To create an iSNS profile, follow these steps:

Procedure

- Step 1** Choose End Devices > iSCSI in the Physical Attributes pane.
You see the iSCSI configuration in the Information pane.
 - Step 2** Select the iSNS tab.
 - Step 3** You see the iSNS profiles configured.
 - Step 4** Click the Create Row icon.
You see the Create iSNS Profiles dialog box.
 - Step 5** Set the ProfileName field to the iSNS profile name that you want to create.
 - Step 6** Set the ProfileAddr field to the IP address of the iSNS server.
 - Step 7** Click **Create** to save these changes.
-

Deleting an iSNS profile

To delete an iSNS profile, follow these steps:

Procedure

- Step 1** Choose End Devices > iSCSI from the Physical Attributes pane.
You see the iSCSI configuration in the Information pane.
 - Step 2** Select the iSNS tab.
You see the iSNS profiles configured.
 - Step 3** Right-click the profile that you want to delete and click the Delete Row icon.
-

Tagging a profile to an interface

To tag a profile to an interface, follow these steps:

Procedure

- Step 1** Choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
You see the Gigabit Ethernet configuration in the Information pane.
 - Step 2** Click the **iSNS** tab.
You see the iSNS profiles configured for these interfaces.
 - Step 3** Set the iSNS ProfileName field to the iSNS profile name that you want to add to this interface.
 - Step 4** Click the **Apply Changes** icon to save these changes.
-

Untagging a profile from an interface

To untag a profile from an interface using Cisco DCNM-SAN, follow these steps:

Procedure

- Step 1** Choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
You see the Gigabit Ethernet Configuration in the Information pane.
 - Step 2** Click the **iSNS** tab.
You see the iSNS profiles configured for these interfaces.
 - Step 3** Right-click the iSNS ProfileName field that you want to untag and delete the text in that field.
 - Step 4** Click the **Apply Changes** icon to save these changes.
-

Configuring iSNS Servers

This section describe how to configure an iSNS server on a Cisco MDS 9000 Family switch.

This section includes the following topics:

Enabling the iSNS Server

Before the iSNS server feature can be enabled, iSCSI must be enabled (see the [Enabling iSCSI, on page 836](#)). When you disable iSCSI, iSNS is automatically disabled. When the iSNS server is enabled on a switch, every IPS port whose corresponding iSCSI interface is up is capable of servicing iSNS registration and query requests from external iSNS clients.

To enable the iSNS server, follow these steps:

Procedure

-
- Step 1** Choose **End Devices > iSNS**.
You see the iSNS configuration in the Information pane.
- Step 2** Click the **Control** tab and select **enable** from the Command drop-down menu for the iSNS server feature.
- Step 3** Click the Apply Changes icon to save this change.
-

What to do next



Note

If you are using VRRP IPv4 addresses for discovering targets from iSNS clients, ensure that the IP address is created using the **secondary** option (see [“Virtual Router Initiation” section on page 41-17](#)).

iSNS Configuration Distribution

You can use the CFS infrastructure to distribute the iSCSI initiator configuration to iSNS servers across the fabric. This allows the iSNS server running on any switch to provide a querying iSNS client a list of iSCSI devices available anywhere on the fabric.

To enable iSNS configuration distribution, follow these steps:

Procedure

-
- Step 1** Choose **End Devices > iSNS**.
You see the iSNS configuration in the Information pane.
- Step 2** Click the **CFS** tab and select **enable** from the Admin drop-down menu for iSNS.
- Step 3** Select **enable** from the Global drop-down menu for iSNS.
- Step 4** Click the **Apply Changes** icon to save this change.
-

Configuring the ESI Retry Count

The iSNS client registers information with its configured iSNS server using an iSNS profile. At registration, the client can indicate an entity status inquiry (ESI) interval of 60 seconds or more. If the client registers with an ESI interval set to zero (0), then the server does not monitor the client using ESI. In such cases, the client's registrations remain valid until explicitly deregistered or the iSNS server feature is disabled.

The ESI retry count is the number of times the iSNS server queries iSNS clients for their entity status. The default ESI retry count is 3. The client sends the server a response to indicate that it is still alive. If the client fails to respond after the configured number of retries, the client is deregistered from the server.

Configuring the Registration Period

The iSNS client specifies the registration period with the iSNS Server. The iSNS Server keeps the registration active until the end of this period. If there are no commands from the iSNS client during this period, then the iSNS server removes the client registration from its database.

If the iSNS client does not specify a registration period, the iSNS server assumes a default value of 0, which keeps the registration active indefinitely. You can also manually configure the registration period on the MDS iSNS Server.

To configure the registration period on an iSNS Server, follow these steps:

Procedure

- Step 1** Choose **End Devices > iSNS**.
You see the iSNS configuration in the Information pane.
 - Step 2** Click the Servers tab.
You see the configured iSNS servers.
 - Step 3** Set the **ESI NonResponse Threshold** field to the ESI retry count value.
 - Step 4** Click the **Apply Changes** icon to save this change.
-

Configuring iSNS Cloud Discovery

This section describes how to configure iSNS cloud discovery and includes the following topics:

Enabling iSNS Cloud Discovery

To enable iSNS cloud discovery, follow these steps:

Procedure

- Step 1** Choose **End Devices > iSNS**.
You see the iSNS configuration in the Information pane.
 - Step 2** Click the **Control** tab and select **enable** from the Command drop-down menu for the cloud discovery feature.
 - Step 3** Click the **Apply Changes** icon to save this change.
-

Initiating On-Demand iSNS Cloud Discovery

To initiate on-demand iSNS cloud discovery, follow these steps:

Procedure

- Step 1** Choose **End Devices > iSNS**.
You see the iSNS configuration in the Information pane.
- Step 2** Click the **Cloud Discovery** tab and check the **Manual Discovery** check box.
- Step 3** Click the **Apply Changes** icon to save this change.
-

Configuring Automatic iSNS Cloud Discovery

To configure automatic iSNS cloud discovery, follow these steps:

Procedure

- Step 1** Choose **End Devices > iSNS**.
You see the iSNS configuration in the Information pane.
- Step 2** Click the **Cloud Discovery** tab and check the **AutoDiscovery** check box.
- Step 3** Click the **Apply Changes** icon to save this change.
-

Configuring iSNS Cloud Discovery Distribution

To configure iSNS cloud discovery CFS distribution, follow these steps:

Procedure

- Step 1** Choose **End Devices > iSNS**.
You see the iSNS configuration in the Information pane.
- Step 2** Click the **CFS** tab and select **enable** from the Admin drop-down menu for the cloud discovery feature.
- Step 3** Select **enable** from the Global drop-down menu for the cloud discovery feature.
- Step 4** Click the **Apply Changes** icon to save this change.
-

Configuring iSNS Cloud Discovery Message Types

You can configure iSNS cloud discovery the type of message to use. By default, iSNS cloud discovery uses ICMP.


Configuration Examples for iSCSI

This section provides three examples of iSCSI virtual target configurations.

Example 1

This example assigns the whole Fibre Channel target as an iSCSI virtual target. All LUNs that are part of the Fibre Channel target are available as part of the iSCSI target (see [Figure 137: Assigning iSCSI Node Names, on page 863](#)).


Figure 137: Assigning iSCSI Node Names



Example 2

This example maps a subset of LUNs of a Fibre Channel target to three iSCSI virtual targets. Each iSCSI target only has one LUN (see [Figure 138: Mapping LUNs to an iSCSI Node Name, on page 863](#)).


Figure 138: Mapping LUNs to an iSCSI Node Name



Example 3

This example maps three subsets of Fibre Channel LUN targets to three iSCSI virtual targets. Two iSCSI targets have one LUN and the third iSCSI target has two LUNs (see [Figure 139: Mapping LUNs to Multiple iSCSI Node Names, on page 864](#)).

Figure 139: Mapping LUNs to Multiple iSCSI Node Names




Example of VSAN Membership for iSCSI Devices

Figure 140: VSAN Membership for iSCSI Interfaces, on page 864 provides an example of VSAN membership for iSCSI devices:

- iSCSI interface 1/1 is a member of VSAN Y.
- iSCSI initiator host A has explicit VSAN membership to VSAN X.
- Three iSCSI initiators (host A, host B, and host C) connect to iSCSI interface 1/1.

Figure 140: VSAN Membership for iSCSI Interfaces




Host A's virtual Fibre Channel N port will be added to VSAN X because of explicit membership for the initiator. The virtual host-B and host-C N ports do not have any explicit membership configuration so they will inherit the iSCSI interface VSAN membership and be part of VSAN Y.

Example of an iSNS Server

The iSNS server provides uniform access control across Fibre Channel and iSCSI devices by utilizing both Fibre Channel zoning information and iSCSI access control information and configuration. An iSCSI initiator acting as an iSNS client only discovers devices it is allowed to access based on both sets of access control information. [Figure 141: Using iSNS Servers in the Cisco MDS Environment, on page 865](#) provides an example of this scenario.

Figure 141: Using iSNS Servers in the Cisco MDS Environment



In [Figure 141: Using iSNS Servers in the Cisco MDS Environment, on page 865](#), iqn.host1 and iqn.host2 are iSCSI initiators. P1 and P2 are Fibre Channel targets. The two initiators are in different zones: Zone 1 consists of iqn.host1 and target P1, and Zone 2 consists of iqn.host2 and target P2. iSNS server functionality is enabled on both switches, SW-1 and SW-2. The registration process proceeds as follows:

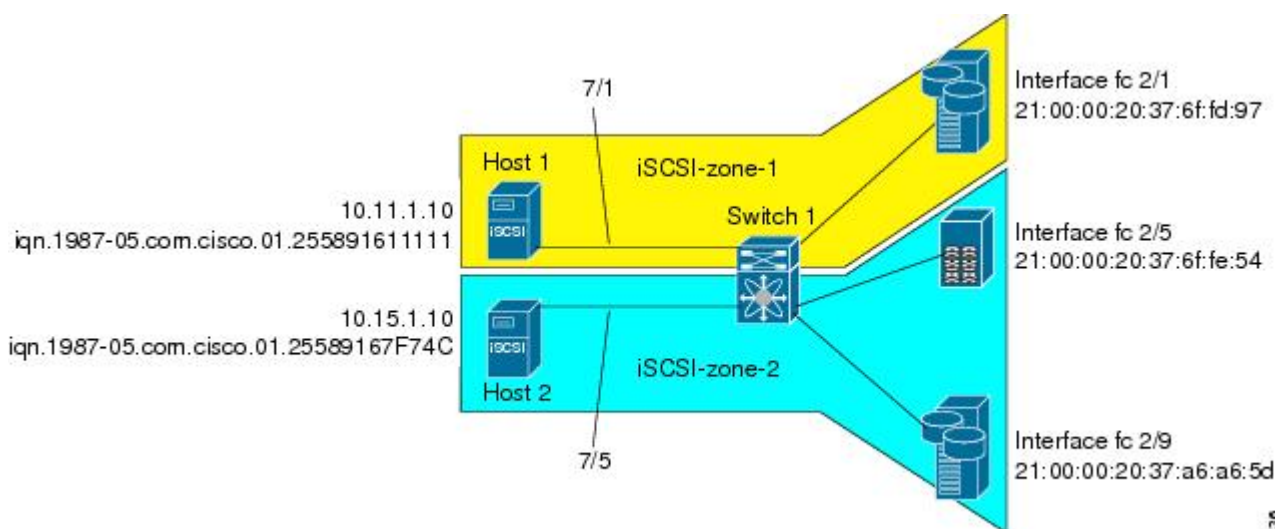
1. Initiator iqn.host1 registers with SW-1, port Gigabitethernet2/1.
2. Initiator iqn.host2 registers with SW-2, port Gigabitethernet3/1.
3. Initiator iqn.host1 issues an iSNS query to SW-1 to determine all accessible targets.
4. The iSNS server in turn queries the Fibre Channel name server (FCNS) to obtain a list of devices that are accessible (that is, in the same zone) by the query originator. This query yields only P1.
5. The iSNS server then queries its own database to convert the Fibre Channel devices to the corresponding iSCSI targets. This is based on the iSCSI configuration, such as virtual-target and its access control setting or whether the dynamic Fibre Channel target import feature is enabled or disabled.
6. The iSNS server sends a response back to the query initiator. This response contains a list all iSCSI portals known to the iSNS server. This means iqn.host1 can choose to log in to target P1 through either SW-1 (at Gigabitethernet 2/1) or SW-2 (at Gigabitethernet 3/1).
7. If the initiator chooses to log in to SW-1 and later that port becomes inaccessible (for example, Gigabitethernet 2/1 goes down), the initiator has the choice to move to connect to target P1 through port Gigabitethernet 3/1 on SW-2 instead.
8. If the target either goes down or is removed from the zone, the iSNS server sends out an iSNS State Change Notification (SCN) message to the initiator so that the initiator can remove the session.

iSCSI Transparent Mode Initiator Example

This examples assumes the following configuration (see [Figure 142: iSCSI Transparent Mode Initiator , on page 866](#)):

- No LUN mapping or LUN masking or any other access control for hosts on the target device
- No iSCSI login authentication (that is, login authentication set to none)
- The topology is as follows:
 - iSCSI interface 7/1 is configured to identify initiators by IP address.
 - iSCSI interface 7/5 is configured to identify initiators by node name.
 - The iSCSI initiator host 1 with IPv4 address 10.11.1.10 and name iqn.1987-05.com.cisco:01.255891611111 connects to IPS port 7/1 is identified using IPv4 address (host 1 = 10.11.1.10).
 - The iSCSI initiator host 2 with IPv4 address 10.15.1.10 and node name iqn.1987-05.com.cisco:01.25589167f74c connects to IPS port 7/5.

Figure 142: iSCSI Transparent Mode Initiator



To configure this example (see [Figure 142: iSCSI Transparent Mode Initiator](#), on page 866), follow these steps:

Procedure

- Step 1** Configure null authentication for all iSCSI hosts in Cisco MDS switches.
 - a) In Cisco DCNM-SAN, choose **End Devices** > **iSCSI** in the Physical Attributes pane.
 - b) Select **none** from the AuthMethod drop-down menu in the Information pane.
 - c) Click the **Apply Changes** icon.
- Step 2** Configure iSCSI to dynamically import all Fibre Channel targets into the iSCSI SAN using auto-generated iSCSI target names.
 - a) In Device Manager, click **IP** > **iSCSI**.
 - b) Click the **Targets** tab.
 - c) Check the **Dynamically Import FC Targets** check box.
 - d) Click **Apply**.
- Step 3** Configure the Gigabit Ethernet interface in slot 7 port 1 with an IPv4 address and enable the interface.

- a) In Cisco DCNM-SAN, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
- b) Select the **IP Address** tab in the Information pane and click **Create Row**.
- c) Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 1.
- d) Click **Create**.
- e) Select the General tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 1.
- f) Click the Apply Changes icon.

Note Host 2 is connected to this port.

Step 4 Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by their IP address, and enable the interface.

- a) In Cisco DCNM-SAN, choose **Switches > Interfaces > FC Logical** in the Physical Attributes pane.
- b) Click the **iSCSI** tab in the Information pane.
- c) Select **ipaddress** from the Initiator ID Mode drop-down menu and click the **Apply Changes** icon.
- d) In Device Manager, choose **Interfaces > Ethernet and iSCSI**.
- e) Click the **iSCSI** tab.
- f) Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 1.
- g) Click **Apply**.

Step 5 Configure the Gigabit Ethernet interface in slot 7 port 5 with an IPv4 address and enable the interface.

- a) In Cisco DCNM-SAN, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
- b) Click the **IP Address** tab in the Information pane and click **Create Row**.
- c) Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 5.
- d) Click **Create**.
- e) Select the General tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 5.
- f) Click the **Apply Changes** icon.

Step 6 Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by node name and enable the interface.

- a) In Cisco DCNM-SAN, choose **Switches > Interfaces > FC Logical** in the Physical Attributes pane.
- b) Click the **iSCSI** tab in the Information pane.
- c) Select **name** from the Initiator ID Mode drop-down menu and click the **Apply Changes** icon.
- d) In Device Manager, choose **Interfaces > Ethernet and iSCSI**.
- e) Click the **iSCSI** tab.
- f) Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 5.
- g) Click **Apply**.

Note Host 1 is connected to this port.

Step 7 Verify the available Fibre Channel targets.

- a) In Device Manager, Choose **FC > Name Server**.
- b) Click the **General** tab.

Step 8 Create a zone named *iscsi-zone-1* with host 1 and one Fibre Channel target in it.

Note Use the IP address of the host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on IP address.

- a) In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database**.

- b) Select VSAN 1 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
- c) Select the **Zones** folder in the left navigation pane and click **Insert**.
- d) Set the Zone Name field to **iscsi-zone-1** and click **OK**.
- e) Select the iscsi-zone-1 folder in the left navigation pane and click **Insert**.
- f) Set the ZoneBy radio button to **WWN**.
- g) Set the Port WWN to the pWWN for the Fibre Channel target (that is, 21:00:00:20:37:6f:fd:97) and click **Add**.
- h) Set the ZoneBy radio button to **iSCSI IP Address/Subnet**.
- i) Set the IP Address/Mask field to the IP Address for Host 1 iSCSI initiator (10.11.1.10) and click **Add**.

Step 9

Create a zone named *iscsi-zone-2* with host 2 and two Fibre Channel targets in it.

Note Use the symbolic node name of the iSCSI host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on node name.

- a) In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database** from the main menu.
- b) Select VSAN 2 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
- c) Select the **Zones** folder in the left navigation pane and click **Insert**.
- d) Set the Zone Name field to **iscsi-zone-2** and click **OK**.
- e) Select the **iscsi-zone-2** folder in the left navigation pane and click **Insert**.
- f) Set the ZoneBy radio button to **WWN**.
- g) Set the Port WWN to the pWWN for one of the Fibre Channel targets (for example, 21:00:00:20:37:6f:fe:5) and click **Add**.
- h) Set the Port WWN to the pWWN for another of the Fibre Channel targets (for example, 21:00:00:20:37:a6:a6:5d) and click **Add**.
- i) Set the ZoneBy radio button to **iSCSI name**.
- j) Set the Port Name field to the symbolic name for host 2 (iqn.1987-05.com.cisco:01.25589167f74c) and click **Add**.

Step 10

Create a zone set, add the two zones as members, and activate the zone set.

Note iSCSI interface is configured to identify all hosts based on node name.

- a) In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database**.
- b) Select VSAN 1 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
- c) Select the **Zoneset** folder in the left navigation pane and click **Insert**.
- d) Set the Zoneset Name to **zonset-iscsi** and click **OK**.
- e) Click on the **zonset-iscsi** folder and click **Insert**.
- f) Set the Zone Name field to **iscsi-zone-1** and click **OK**.
- g) Set the Zone Name field to **iscsi-zone-2** and click **OK**.
- h) Click **Activate** to activate the new zone set.
- i) Click **Continue Activation** to finish the activation.

Step 11

Bring up the iSCSI hosts (host 1 and host 2).

Step 12

Show all the iSCSI sessions.

- a) In Device Manager, choose **Interfaces > Monitor > Ethernet**.
- b) Click the **iSCSI connections** tab to show all the iSCSI sessions.
- c) In Device Manager, choose **IP > iSCSI** and select the **Session Initiators** tab.
- d) Click **Details**.


- Step 13** In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane to verify the details of the two iSCSI initiators
- Step 14** In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database** to view the active zone set. The iSCSI initiators' FC IDs are resolved.
- Step 15** In Device Manager, Choose **FC > Name Server**. The Fibre Channel name server shows the virtual N ports created for the iSCSI hosts.
- Step 16** In Device Manager, Choose **FC > Name Server**.
- Step 17** Click the **Advanced** tab. Verify the detailed output of the iSCSI initiator nodes in the Fibre Channel name server.

Target Storage Device Requiring LUN Mapping Example

This example scenario 2 assumes the following configuration (see [Figure 143: Target Storage Device with LUN Mapping, on page 869](#)):

- Access control is based on Fibre Channel zoning.
- There is target-based LUN mapping or LUN masking.
- There is no iSCSI authentication (none).
- The iSCSI initiator is assigned to different VSANs.

Figure 143: Target Storage Device with LUN Mapping



To configure this example (see [Figure 143: Target Storage Device with LUN Mapping, on page 869](#)), follow these steps:

Procedure

- Step 1** Configure null authentication for all iSCSI hosts.
- a) In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane.

- b) Select **none** from the AuthMethod drop-down menu in the Information pane.
- c) Click the **Apply Changes** icon.

Step 2 Configure iSCSI to dynamically import all Fibre Channel targets into the iSCSI SAN using auto-generated iSCSI target names.

- a) In Device Manager, click **IP > iSCSI**.
- b) Click the **Targets** tab.
- c) Check the **Dynamically Import FC Targets** check box.
- d) Click **Apply**.

Step 3 Configure the Gigabit Ethernet interface in slot 7 port 1 with an IPv4 address and enable the interface.

- a) In Cisco DCNM-SAN, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
- b) Select the **IP Address** tab in the Information pane and click **Create Row**.
- c) Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 1.
- d) Click **Create**.
- e) Click the **General** tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 1.
- f) Click the **Apply Changes** icon.

Step 4 Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by their IP address and enable the interface.

- a) In Cisco DCNM-SAN, choose **Switches > Interfaces > FC Logical** in the Physical Attributes pane.
- b) Select the **iSCSI** tab in the Information pane.
- c) Select **ipaddress** from the Initiator ID Mode drop-down menu and click the **Apply Changes** icon.
- d) In Device Manager, choose **Interfaces > Ethernet and iSCSI**.
- e) Click the **iSCSI** tab.
- f) Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 1.
- g) Click **Apply**.

Step 5 Configure the Gigabit Ethernet interface in slot 7 port 5 with the IPv4 address and enable the interface.

- a) In Cisco DCNM-SAN, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
- b) Click the **IP Address** tab in the Information pane and click **Create Row**.
- c) Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 5.
- d) Click **Create**.
- e) Select the **General** tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 5.
- f) Click the **Apply Changes** icon.

Step 6 Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by IP address and enable the interface.

- a) In Cisco DCNM-SAN, choose **Switches > Interfaces > FC Logical** in the Physical Attributes pane.
- b) Click the **iSCSI** tab in the Information pane.
- c) Select **ipaddress** from the Initiator ID Mode drop-down menu and click the **Apply Changes** icon.
- d) In Device Manager, choose **Interfaces > Ethernet and iSCSI**.
- e) Click the **iSCSI** tab.
- f) Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 5.
- g) Click **Apply**.

Step 7 Configure for static pWWN and nWWN for host 1.

- a) In Device Manager, choose **IP > iSCSI**.
- b) Click the **Initiators** tab.
- c) Check the **Node Address Persistent** and **Node Address System-assigned** check boxes the Host 1 iSCSI initiator.
- d) Click **Apply**.

Step 8 Configure for static pWWN for Host 2.

- a) In Device Manager, Choose **IP > iSCSI**.
- b) Click the **Initiators** tab.
- c) Right-click on the Host 2 iSCSI initiator and click Edit pWWN.
- d) Select 1 from the System-assigned Num field and click **Apply**.

Step 9 View the configured WWNs.

Note The WWNs are assigned by the system. The initiators are members of different VSANs.

- a) In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane.
- b) Click the **Initiators** tab.

Step 10 Create a zone for Host 1 and the iSCSI target in VSAN 1.

Note Use the IP address of the host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on IP address.

- a) In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database**.
- b) Select VSAN 1 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
- c) Select the **Zones** folder in the left navigation pane and click **Insert**.
- d) Set the Zone Name field to **iscsi-zone-1** and click **OK**.
- e) Select the iscsi-zone-1 folder in the left navigation pane and click **Insert**.
- f) Set the ZoneBy radio button to **WWN**.
- g) Set the Port WWN to the pWWN for the Fibre Channel target (that is, 21:00:00:20:37:6f:fd:97). and click **Add**.
- h) Set the ZoneBy radio button to **iSCSI IP Address/Subnet**.
- i) Set the IP Address/Mask field to the IP Address for Host 1 iSCSI initiator (10.11.1.10) and click **Add**.

Note Fibre Channel storage for zone membership for the iSCSI initiator, either the iSCSI symbolic node name or the pWWN, can be used. In this case, the pWWN is persistent.

Step 11 Create a zone set in VSAN 1 and activate it.

- a) In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database**.
- b) Select VSAN 1 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
- c) Select the **Zoneset** folder in the left navigation pane and click **Insert**.
- d) Set the Zoneset Name to **zonset-iscsi-1** and click **OK**.
- e) Click on the **zonset-iscsi-1** folder and click **Insert**.
- f) Set the Zone Name field to **iscsi-zone-1** and click **OK**.
- g) Click **Activate** to activate the new zone set.
- h) Click **Continue Activation** to finish the activation.

Step 12 Create a zone with host 2 and two Fibre Channel targets.

Note If the host is in VSAN 2, the Fibre Channel targets and zone must also be in VSAN 2.

Note iSCSI interface is configured to identify all hosts based on node name.

- a) In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database**.
- b) Select VSAN 2 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
- c) Select the **Zones** folder in the left navigation pane and click **Insert**.
- d) Set the Zone Name field to **iscsi-zone-2** and click **OK**.
- e) Select the **iscsi-zone-2** folder in the left navigation pane and click **Insert**.
- f) Set the ZoneBy radio button to **WWN**.
- g) Set the Port WWN to the pWWN for one of the Fibre Channel targets (for example, 21:00:00:20:37:6f:fe:5) and click **Add**.
- h) Set the Port WWN to the pWWN for another of the Fibre Channel targets (for example, 21:00:00:20:37:a6:a6:5d) and click **Add**.
- i) Set the ZoneBy radio button to **iSCSI IP Address/Subnet**.
- j) Set the IP Address/Mask field to the IP Address for Host 2 iSCSI initiator (10.15.1.11) and click **Add**.

Step 13 Create a zone set in VSAN 2 and activate it.

- a) In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database**.
- b) Select VSAN 2 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
- c) Select the **Zoneset** folder in the left navigation pane and click **Insert**.
- d) Set the Zoneset Name to **zonset-iscsi-2** and click **OK**.
- e) Click on the **zoneset-iscsi-2** folder and click **Insert**.
- f) Set the Zone Name field to **iscsi-zone-2** and click **OK**.
- g) Click **Activate** to activate the new zone set.
- h) Click **Continue Activation** to finish the activation.

Step 14 Start the iSCSI clients on both hosts.

Step 15 Show all the iSCSI sessions.

- a) In Device Manager, choose **Interface > Monitor > Ethernet** and select the **iSCSI connections** tab to show all the iSCSI sessions.
- b) In Device Manager, choose **IP > iSCSI** and select the **Session Initiators** tab.
- c) Click **Details**.

Step 16 In Cisco DCNM- SAN, choose **End Devices > iSCSI** in the Physical Attributes pane to verify the details of the two iSCSI initiators.

Step 17 In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database** to view the active zone set. The iSCSI initiators' FC IDs are resolved.

Step 18 In Device Manager, choose **FC > Name Server**. The Fibre Channel name server shows the virtual N ports created for the iSCSI hosts.

Step 19 In Device Manager, Choose **FC > Name Server**.

Step 20 Click the **Advanced** tab. Verify the detailed output of the iSCSI initiator nodes in the Fibre Channel name server.

Field Descriptions for iSCSI

The following are the field descriptions for iSCSI.

Ethernet Interfaces iSCSI

Field	Description
Description	An alias name for the interface as specified by a network manager.
Speed	Operational speed.
PhysAddress	The interface's WWN.
Admin	The desired state of the interface.
Oper	The current operational state of the interface.
LastChange	When the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this value contains a N/A value.
PortVSAN	The VSAN that the interface belongs to.
ForwardingMode	Use Store and Forward if the HBA has problems with passthrough.
Initiator ID Mode	How the initiator is identified on this interface, either by its iSCSI name (name) or by its IP address (ipaddress).
Enable	The initiator proxy mode for this interface. If true, then all the initiators coming on this interface would use the initiator configuration provided by this interface. The initiator configuration include port WWN and node WWN.
Assignment	How the initiator proxy mode FC addresses are assigned. If auto, then the FC addresses are automatically assigned. If it is manual, then they have to be manually configured.
Port WWN	The Port FC address used by the initiators on this interface when the initiator proxy mode is on.
Node WWN	The Node FC address used by the initiators on this interface when the initiator proxy mode is on.

Ethernet Interfaces iSCSI TCP

Field	Description
Local Port	Local interface TCP port.
SACK	Indicates if the Selective Acknowledgement (SACK) option is enabled or not.
KeepAlive	The TCP keepalive timeout for this iSCSI interface. If the value is 0, the keepalive timeout feature is disabled.
MinTimeout	The TCP minimum retransmit time.
Max	The TCP maximum retransmissions.
SendBufferSize	The TCP send buffer size.

Field	Description
MinBandwidth	The TCP minimum bandwidth.
MaxBandwidth	The TCP maximum bandwidth.
Estimated Round Trip	The estimated round trip delay of network pipe used for B-D product computation. The switch can use this to derive the TCP window to advertise.
QoS	The TCP QoS code point.
PMTU Enable	Indicates if the Path MTU discovery option is enabled or not.
PMTU Reset Timeout	The PMTU reset timeout.
Connections Normal	The number of normal iSCSI connections.
Connections Discovered	The number of discovery iSCSI connections.
CWM Enable	If true, congestion window monitoring is enabled. If false, it is disabled.
CWM Burst Size	The maximum burst sent after a TCP sender idle period.
Max Jitter	The maximum delay variation (not due to congestion) that can be experienced by TCP connections on this interface.
Port	The local TCP port of this interface.

Ethernet Interface Monitor iSCSI Connections

Field	Description
RxBytes	Total number of bytes received on an iSCSI session.
TxBytes	Total number of bytes transmitted on an iSCSI session.
IPSEC	A collection of objects for iSCSI connection statistics.

iSCSI Connection

Field	Description
LocalAddr	The local Internet network address used by this connection.
RemoteAddr	The remote Internet network address used by this connection.
CID	The iSCSI Connection ID for this connection.

Field	Description
State	The current state of this connection, from an iSCSI negotiation point of view. <ul style="list-style-type: none"> • login- The transport protocol connection has been established, but a valid iSCSI login response with the final bit set has not been sent or received. • full- A valid iSCSI login response with the final bit set has been sent or received. • logout- A valid iSCSI logout command has been sent or received, but the transport protocol connection has not yet been closed.
MaxRecvDSLen	The maximum data payload size supported for command or data PDUs in use within this connection. Note that the size of reported in bytes even though the negotiation is in 512 k blocks.
SendMarker	Indicates whether or not this connection is inserting markers in its outgoing data stream.
HeaderDigest	The iSCSI header digest scheme in use within this connection.
DataDigest	The iSCSI data digest scheme in use within this connection.

iSCSI Initiators

Field	Description
Name or IP Address	A character string that is a globally unique identifier for the node represented by this entry.
VSAN Membership	The list of configured VSANs the node represented by this entry can access.
Dynamic	If true, then the node represented by this entry is automatically discovered.
Initiator Type	Indicates whether the node is a host that participates in iSCSI load balancing.
Persistent Node WWN	If true, then the same FC address is assigned to the node if it were to be represented again in the FC domain with the same node name. Note that the node FC address is either automatically assigned or manually configured.
SystemAssigned Node WWNN	If true, the FC address is automatically assigned to this node. If false, then the FC address has to be configured manually.
Node WWN	The persistent FC address of the node.
Persistent Port WWN	If true, then the same FC address is assigned to the ports of the node if it were to be represented again in the FC domain with the same node name.
Port WWN	All the FC port addresses associated with this node.
AuthUser	This is the only CHAP user name that the initiator is allowed to log in with.
Target UserName	(Optional) The user name to be used for login. If you do not supply a username, the global user name is used.
Target Password	(Optional) The password to be used for login. If you do not supply a password, the global password is used.

Field	Description
Load Metric	A configured load metric of this iSCSI initiator for the purpose of iSCSI load balancing.
Auto Zone Name	The zone name that is used when the system creates automatic zone for this initiator's specific list of targets.

iSCSI Targets

Field	Description
Dynamically Import FC Targets	Check this option to dynamically import FC targets into the iSCSI domain. A target is not imported if it already exists in the iSCSI domain.
iSCSI Name	The iSCSI name of the node represented by this entry.
Dynamic	Indicates if the node represented by this entry was either automatically discovered or configured manually.
Primary Port WWN	The FC address for this target.
Secondary Port WWN	The optional secondary FC address for this target. This is the FC address used if the primary cannot be reached.
LUN Map iSCSI	The configured default logical unit number of this LU.
LUN Map FC Primary	The logical unit number of the remote LU for the primary port address.
LUN Map FC Secondary	The logical unit number of the remote LU for the secondary port address.
Initiator Access All	If true, then all the initiators can access this target even those which are not in the initiator permit list of this target. If false, then only initiators which are in the permit list are allowed access to this target.
Initiator Access List	Lists all the iSCSI nodes that are permitted to access the node represented by this entry. If AllAllowed is false and the value of List is empty, then no initiators are allowed to access this target.
Advertised Interfaces	Lists all the interfaces on which the target could be advertised.
Trespass Mode	The trespass mode for this node. Every iSCSI target represents one or more port(s) on the FC target. If true, the node instructs the FC node to present all LUN I/O requests to secondary port if the primary port is down.
RevertToPrimaryPort	Indicates if it is required to revert back to primary port if the FC target comes back online.

iSCSI Session Initiators

Field	Description
Name or IP Address	The name or IP address of the initiator port.

Field	Description
Alias	The initiator alias acquired at login.

iSCSI Global

Field	Description
AuthMethod	The authentication method.
InitiatorIdleTimeout	The time for which the gateway (representing a FC target) waits from the time of last iSCSI session to a iSCSI initiator went down, before purging the information about that iSCSI initiator.
iSLB ZonesetActivate	Checking this option performs automatic zoning associated with the initiator targets
DynamicInitiator	This field determines how dynamic iSCSI initiators are created. Selecting the iSCSI option (default) creates dynamic iSCSI initiators. If you select iSLB then the an iSLB dynamic initiator is created. Selecting the deny option does not allow dynamic creation of the initiators.
Target UserName	The default user name used for login. If an initiator user name is specified, that user name is used instead.
Target Password	The default password used for login. If an initiator password is specified, that password is used instead.

iSCSI Session Statistics

Field	Description
PDU Command	The count of command PDUs transferred on this session.
PDU Response	The count of response PDUs transferred on this session.
Data Tx	The count of data bytes that were transmitted by the local iSCSI node on this session.
Data Rx	The count of data bytes that were received by the local iSCSI node on this session.
Errors Digest	Authentication errors.
Errors CxnTimeout	Connection timeouts.

iSCSI iSLB VRRP

Field	Description
VrId, IpVersion	The virtual router number and the IP version (IPv4, IPv6, or DNS).
Load Balance	Indicates whether load balancing is enabled.

iSCSI Initiator Access

Field	Description
Initiator Name	The iSCSI node name.

iSCSI Initiator PWWN

Field	Description
Port WWN	The FC address for this entry.

iSCSI Sessions

Field	Description
Type	Type of iSCSI session: <ul style="list-style-type: none"> • normal—session is a normal iSCSI session • discovery—session is being used only for discovery.
TargetName	If Direction is Outbound, this will contain the name of the remote target.
Vsan ID	The VSAN to which this session belongs to.
ISID	The initiator-defined portion of the iSCSI session ID.
TSIH	The target-defined identification handle for this session.

iSCSI Sessions Detail

Field	Description
ConnectionNumber	The number of transport protocol connections that currently belong to this session.
ImmediateData	Whether the initiator and target have agreed to support immediate data on this session.
Initial	If true, the initiator must wait for a Ready-To-Transfer before sending to the target. If false, the initiator may send data immediately, within limits set by FirstBurstSize and the expected data transfer length of the request.
MaxOutstanding	The maximum number of outstanding Ready-To-Transfers per task within this session.
First	The maximum length supported for unsolicited data sent within this session.
Max	The maximum number of bytes which can be sent within a single sequence of Data-In or Data-Out PDUs.

Field	Description
Sequence	If false, indicates that iSCSI data PDU sequences may be transferred in any order. If true indicates that data PDU sequences must be transferred using continuously increasing offsets, except during error recovery.
PDU	If false, iSCSI data PDUs within sequences may be in any order. If true indicates that data PDUs within sequences must be at continuously increasing addresses, with no gaps or overlay between PDUs.

iSNS Details iSCSI Nodes

Field	Description
Name	The iSCSI name of the initiator or target associated with the storage node.
Type	The Node Type bit-map defining the functions of this iSCSI node, where 31 is a Target, 30 is an Initiator, 29 is a Control, and all others are reserved.
Alias	The Alias name of the iSCSI node.
ScnBitmap	The State Change Notification (SCN) bitmap for a node.
WWN Token	An optional globally unique 64-bit integer value that can be used to represent the world wide node name of the iSCSI device in a Fibre Channel fabric.
AuthMethod	The iSCSI authentication method enabled for this iSCSI Node.

iSCSI User

Field	Description
iSCSI User	The name of the iSCSI user.
Password	The password of the iSCSI user.

Edit iSCSI Advertised Interfaces

Field	Description
Num	The number of the iSCSI target.
Interface	The interface over which the target is to be advertised.



CHAPTER 41

Configuring IP Services

- [Configuring IP Services, on page 881](#)

Configuring IP Services

This chapter includes the following topics:

Information About IP Services

Cisco MDS 9000 Family switches can route IP traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature is used to route traffic between VSANs. To do so, each VSAN must be in a different IP subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMSs):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding on in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

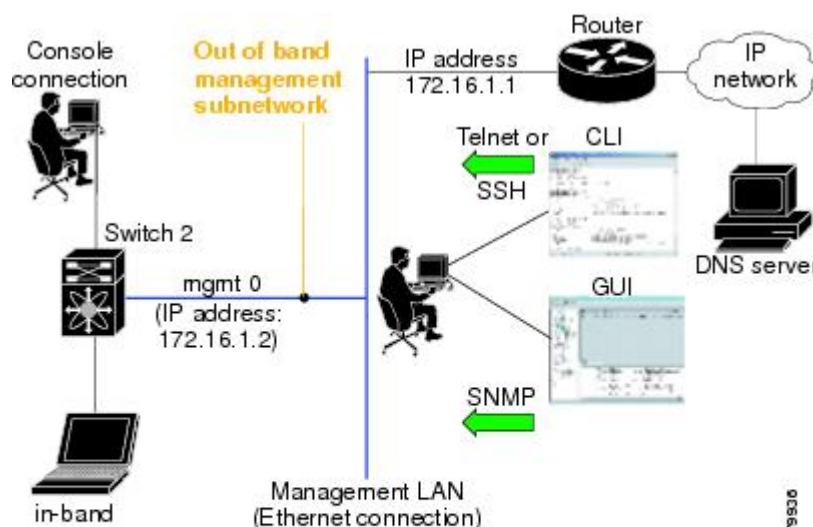
Switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. VRRP is a restartable application that provides a redundant, alternate path to the gateway switch.

This section includes the following topics:

Traffic Management Services

In-band options are compliant with and use the RFC 2625 standards. An NMS host running the IP protocol over an Fibre Channel interface can access the switch using the IPFC functionality. If the NMS does not have a Fibre Channel HBA, in-band management can still be performed using one of the switches as an access point to the fabric as shown in [Figure 144: Management Access to Switches, on page 882](#).

Figure 144: Management Access to Switches



Management Interface Configuration

The management interface on the switch allows multiple simultaneous Telnet or SNMP sessions. You can remotely configure the switch through the management interface, but first you must configure IP version 4 (IPv4) parameters (IP address, subnet mask) or an IP version 6 (IPv6) address and prefix length so that the switch is reachable.

On director class switches, a single IP address is used to manage the switch. The active supervisor module's management (mgmt0) interface uses this IP address. The mgmt0 interface on the standby supervisor module remains in an inactive state and cannot be accessed until a switchover happens. After a switchover, the mgmt0 interface on the standby supervisor module becomes active and assumes the same IP address as the previously active supervisor module.



Note

The port on the Ethernet switch to which the MDS management interface is connected should be configured as a host port (also known as access port) instead of a switch port. Spanning tree configuration for that port (on the Ethernet switch) should be disabled. This helps avoid the delay in the MDS management port coming up due to delay from Ethernet spanning tree processing that the Ethernet switch would run if enabled.



Note

Before you begin to configure the management interface manually, obtain the switch's IP address and IP subnet mask. Also make sure the console cable is connected to the console port.

About the Default Gateway

You can configure a default gateway IPv4 address on your Cisco MDS 9000 Family switch.

The default gateway IPv4 address should be configured along with the IPv4 static routing attributes (IP default network, destination prefix, and destination mask, and next hop address). If you configure the static route IP forwarding and the default-network details, these IPv4 addresses will be used regardless of the default-gateway being enabled or disabled.

The default gateway IPv4 address should be configured along with the IPv4 static routing attributes commands (IP default network, destination prefix, and destination mask, and next hop address).



Tip If you configure the static route IP forwarding and the default-network details, these IPv4 addresses will be used regardless of the default-gateway being enabled or disabled. If these IP addresses are configured but not available, the switch will fall back to using the default gateway IP address, if you have configured it. Be sure to configure IP addresses for all entries in the switch.

IPv4 Default Network Configuration


If you assign the IPv4 default network address, the switch considers routes to that network as the last resort. If the IPv4 default network address is not available, the switch uses the IPv4 default gateway address. For every network configured with the IPv4 default network address, the switch flags that route as a candidate default route, if the route is available.



Tip If you configure the static route IP forwarding and the default network details, these IPv4 addresses will be used regardless of the default gateway being enabled or disabled. If these IPv4 addresses are configured and not available, the switch will fall back to using the default gateway IPv4 address, if you have configured it. Be sure to configure IPv4 addresses for all entries in the switch if you are using IPv4.

When the Ethernet interface is configured, the switch should point to the gateway router for the IP network. The host accesses the gateway using a gateway switch. This gateway switch is configured as the default gateway. The other switches in the fabric that are connected to the same VSAN as the gateway switch can also be connected through the gateway switch. Every interface connected to this VSAN should be configured with the VSAN IPv4 address of the gateway switch as shown in [Figure 145: Overlay VSAN Functionality, on page 884](#)

Figure 145: Overlay VSAN Functionality



In the above figure, switch A has the IPv4 address 1.12.11.1, switch B has the IPv4 address 1.12.11.2, switch C has the IPv4 address 1.12.11.3, and switch D has the IPv4 address 1.12.11.4. Switch A is the gateway switch with the Ethernet connection. The NMS uses the IPv4 address 1.1.1.10 to connect to the gateway switch. Frames forwarded to any switch in the overlaid VSAN 1 are routed through the gateway switch. Configuring the gateway switch's IPv4 address (1.12.11.1) in the other switches enable the gateway switch to forward the frame to the intended destination. Similarly, if a non-gateway switch in the VSAN forwards a frame to the Ethernet, the frame is routed through the gateway switch.

When forwarding is disabled (default), IP frames are not sent from one interface to another. In these cases, the software performs local IP routing between two switches using the in-band option for Fibre Channel traffic and the mgmt0 option for Ethernet traffic.

When a VSAN is created, a VSAN interface is not created automatically. You need to specifically create the interface.

IPFC

IPFC provides IP forwarding on in-band switch management over a Fibre Channel interface (rather than out-of-band using the Gigabit Ethernet mgmt 0 interface). You can use IPFC to specify that IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.

Once the VSAN interface is created, you can specify the IP address for that VSAN. You can assign an IPv4 address or an IPv6 address.

About IPv4 Static Routes

Static routing is a mechanism to configure IPv4 routes on the switch. You can configure more than one static route.

If a VSAN has multiple exit points, configure static routes to direct traffic to the appropriate gateway switch. IPv4 routing is disabled by default on any gateway switch between the out-of-band management interface and the default VSAN, or between directly connected VSANs.

If your network configuration does not need an external router, you can configure IPv4 static routing on your MDS switch.

About Overlay VSANs

VSANs enable deployment of larger SANs by overlaying multiple logical SANs, each running its own instance of fabric services, on a single large physical network. This partitioning of fabric services reduces network instability by containing fabric reconfiguration and error conditions within an individual VSAN. VSANs also provide the same isolation between individual VSANs as physically separated SANs. Traffic cannot cross VSAN boundaries and devices may not reside in more than one VSAN. Because each VSAN runs separate instances of fabric services, each VSAN has its own zone server and can be zoned in exactly the same way as SANs without VSAN capability.

About VRRP

Cisco MDS 9000 Family switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. VRRP provides a redundant alternative path to the gateway switch, which has connectivity to the NMS. VRRP has the following features:


- VRRP is a restartable application.
- When a VRRP master fails, the VRRP backup takes over within three times the advertisement time.
- VRRP over Ethernet, VRRP over VSAN, and Fibre Channel functions are implemented as defined in RFC 2338 and the draft-ietf-vrrp-ipv6 specification.
- A virtual router is mapped to each VSAN and Ethernet interface with its unique virtual router IP, virtual router MAC, and VR ID.
- VR IDs can be reused in multiple VSANs with different virtual router IP mapping.
- Both IPv4 and IPv6 is supported.
- The management interface (mgmt 0) supports only one virtual router group. All other interfaces each support up to seven virtual router groups, including both IPv4 and IPv6 combined. Up to 255 virtual router groups can be assigned in each VSAN.
- VRRP security provides three options, including no authentication, simple text authentication, and MD5 authentication.



Note If you are using IPv6, you must either configure an IPv6 address on the interface or enable IPv6 on the interface.

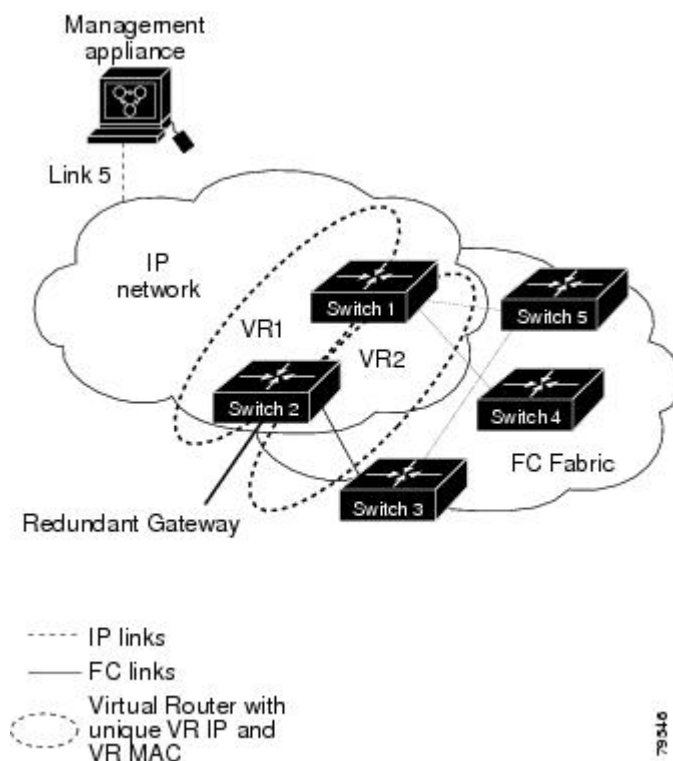
In [Figure 146: VRRP Functionality, on page 886](#), switch A is the VRRP master and switch B is the VRRP backup switch. Both switches have an IP address to VRRP mapping configured. The other switches set switch A as the default gateway. If switch A fails, the other switches do not have to change the routing configurations as switch B automatically becomes the master and takes over the function of a gateway.

Figure 146: VRRP Functionality



In [Figure 147: Redundant Gateway](#), on page 886, the fabric example has two virtual router groups (VR1 and VR2) because a virtual router cannot span across different types of interfaces. In both switch 1 and switch 2, the Ethernet interface is in VR1 and the FC interface is in VR2. Each virtual router is uniquely identified by the VSAN interface and the VR ID.

Figure 147: Redundant Gateway



DNS Server Configuration

The DNS client on the switch communicates with the DNS server to perform the IP address-name server correspondence.

The DNS server may be dropped after two attempts because of one of the following reasons:

- The IP address or the switch name is wrongly configured.
- The DNS server is not reachable because external reasons (reasons beyond our control).



Note When accessing a Telnet host, if the DNS server is not reachable (for any reason) the switch login prompt may take a longer time to appear. If so, verify that the DNS server is accurately configured and reachable.

Guidelines and Limitations

Follow these guidelines to configure IPFC:

1. Create the VSAN to use for in-band management, if necessary.
2. Configure an IPv4 address and subnet mask for the VSAN interface.
3. Enable IPv4 routing.
4. Verify connectivity.

Default Settings

[Table 108: Default DNS Settings](#) , on page 887 lists the default settings for DNS features.

Table 108: Default DNS Settings

Parameters	Default
Domain lookup	Disabled
Domain name	Disabled
Domains	None
Domain server	None
Maximum domain servers	6

[Table 109: Default VRRP Settings](#) , on page 887 lists the default settings for VRRP features.

Table 109: Default VRRP Settings

Parameters	Default
Virtual router state	Disabled
Maximum groups per VSAN	255
Maximum groups per Gigabit Ethernet port	7

Parameters	Default
Priority preemption	Disabled
Virtual router priority	100 for switch with secondary IP addresses 255 for switches with the primary IP address
Priority interface state tracking	Disabled
Advertisement interval	1 second for IPv4 100 centiseconds for IPv6

Configuring IP Services

This section includes the following topics:

Configuring Management Interface

To configure the mgmt0 Ethernet interface using Device Manager for IPv6, follow these steps:

Procedure

-
- Step 1** Select **Interface > Mgmt > Mgmt0**.
 - Step 2** Enter the description.
 - Step 3** Select the administrative state of the interface.
 - Step 4** Check the **CDP** check box to enable CDP.
 - Step 5** Enter the IP address mask.
 - Step 6** Click **Apply** to apply the changes.
-

Configuring the Default Gateway

To configure an IP route, follow these steps:

Procedure

-
- Step 1** Select **Switches > Interfaces > Management**, and select **IP** in the Physical Attributes pane.
 - Step 2** Click the **Route** tab in the information pane.

You see the IP route window showing the switch name, destination, mask, gateway, metric, interface, and active status of each IP route.
 - Step 3** Click the **Create Row** icon to add a new IP route.
 - Step 4** Complete the fields in this window.
 - Enter the switch name in the **Switch** field.

- Configure a static route, by entering the destination network ID and subnet mask in the Routedest and Mask fields.
- Configure a default gateway by entering the IP address of the seed switch in the Gateway field.
- Set the Metric and Interface fields.

Note With Cisco NX-OS Release 4.2(1) and later, CPP interfaces also are available for selection when you create a new IP route.

Step 5 Click the Create icon.

Configuring an IP route or identify the default gateway using Device Manager

To configure an IP route or identify the default gateway using Device Manager, follow these steps:

Procedure

Step 1 Choose IP > Routes.

You see the IP Routes window.

Step 2 Create a new IP route or identify the default gateway on a switch by clicking Create.

Step 3 Complete the fields in this window.

- Enter the switch name in the Switch field.
- Configure a static route, by entering the destination network ID and subnet mask in the Routedest and Mask fields.
- Configure a default gateway by entering the IP address of the seed switch in the Gateway field.
- Set the Metric and Interface fields.

Note With Cisco NX-OS Release 4.2(1) and later, CPP interfaces also are available for selection when you create a new IP route.

If you choose the CPP interface, the switch uses the input CPP-assigned IP address and mask to generate the IP route prefix.

Step 4 Click Create to add the IP route.

Note You cannot delete the switch-generated IP route for the CPP interface. If you try to delete the IP route for the CPP interface, SNMP displays this error message: ip: route type not supported.

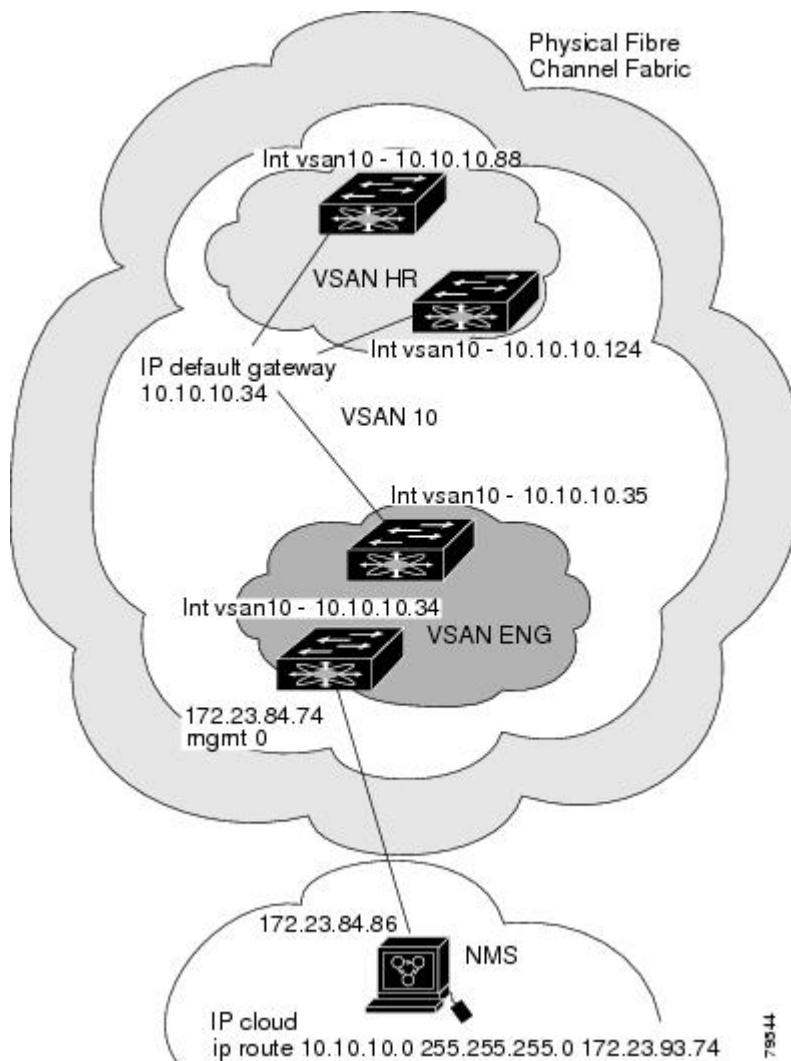
Configuring Overlay VSANs

To configure an overlay VSAN, follow these steps:

Procedure

- Step 1** Add the VSAN to the VSAN database on all switches in the fabric.
- Step 2** Create a VSAN interface for the VSAN on all switches in the fabric. Any VSAN interface belonging to the VSAN has an IP address in the same subnet. Create a route to the IPFC cloud on the IP side.
- Step 3** Configure a default route on every switch in the Fibre Channel fabric pointing to the switch that provides NMS access.
- Step 4** Configure the default gateway (route) and the IPv4 address on switches that point to the NMS as shown in [Figure 148: Overlay VSAN Configuration Example](#) , on page 890.

Figure 148: Overlay VSAN Configuration Example



What to do next



Note To configure the management interface displayed in [Figure 148: Overlay VSAN Configuration Example](#), on [page 890](#), set the default gateway to an IPv4 address on the Ethernet network.

The following procedure configures an overlay VSAN in one switch. This procedure must be repeated for each switch in the fabric.

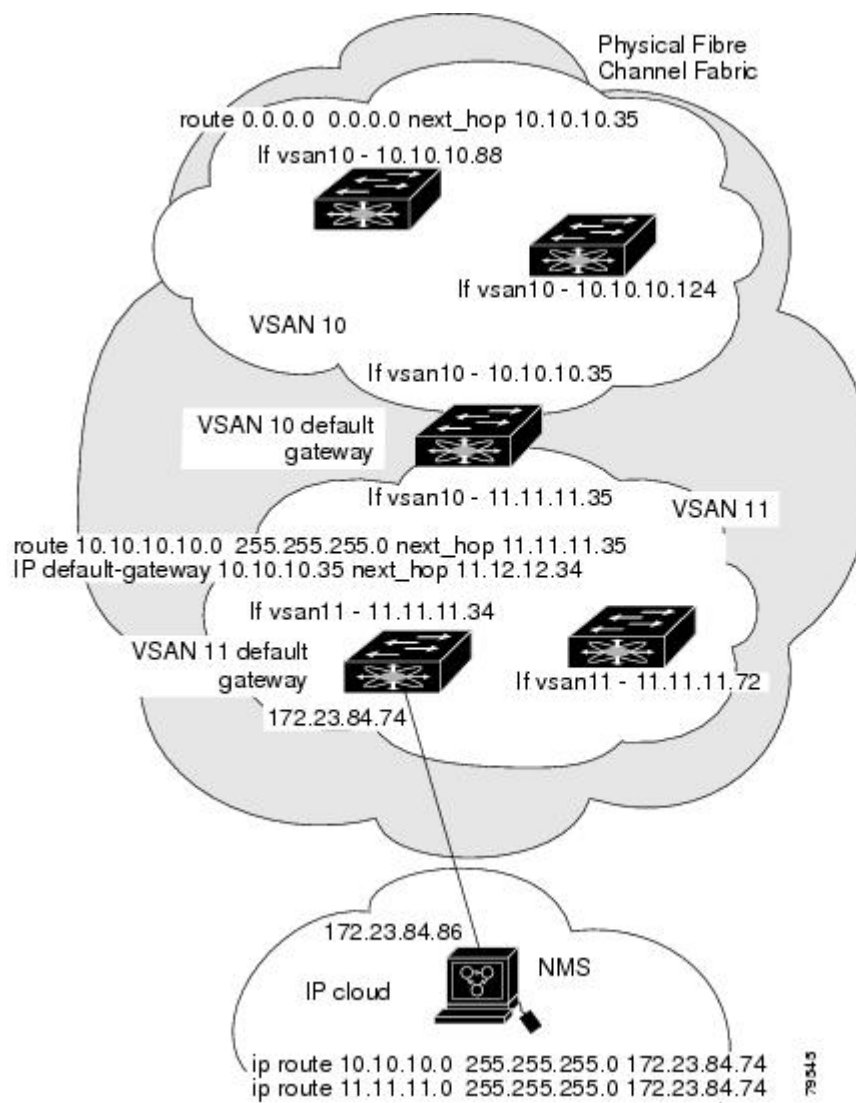
Configuring Multiple VSANs

More than one VSAN can be used to segment the management network in multiple subnets. An active interface must be present on the switch for the VSAN interface to be enabled.

To configure multiple VSANs, follow these steps:

Procedure

- Step 1** Add the VSAN to the VSAN database on any switch in the fabric.
- Step 2** Create a VSAN interface for the appropriate VSAN on any switch in the fabric.
- Step 3** Assign an IP address on every VSAN interface on the same subnet as the corresponding VSAN.
- Step 4** Define the multiple static routes on the Fibre Channel switches and the IP cloud as shown in [Figure 41-6](#).



Configuring VRRP

This section describes how to configure VRRP and includes the following topics:

Adding and Deleting a Virtual Router

All VRRP configurations should be replicated across switches in a fabric that runs VRRP.



Note

The total number of VRRP groups that you can configure on a Gigabit Ethernet port, including main interfaces and subinterfaces, cannot exceed seven. This limitation applies to both IPv4 and IPv6 groups.

Virtual Router Initiation

By default, a virtual router is always disabled. VRRP can be configured only if this state is enabled. Be sure to configure at least one IP address, either IPv4 or IPv6, before attempting to enable a VR.

Adding Virtual Router IP Addresses

One virtual router IP address can be configured for a virtual router. If the configured IP address is the same as the interface IP address, this switch automatically owns the IP address. You can configure either an IPv4 address or an IPv6 address.

According to the VRRP specification, the master VRRP router drops the packets addressed to the virtual router's IP address because the virtual router is only intended as a next-hop router to forward packets. In MDS switches however, some applications require that packets addressed to virtual router's IP address be accepted and delivered to them. By using the **secondary** option to the virtual router IPv4 address, the VRRP router will accept these packets when it is the master.

To manage IP addresses for virtual routers from Device Manager, follow these steps:

Procedure

-
- | | |
|---------------|--------------------------------------------------------------------------------------------|
| Step 1 | Choose IP > VRRP. You see the Operations tab of the VRRP dialog box. |
| Step 2 | Click the IP Addresses tab on the VRRP dialog box. |
| Step 3 | To create a new VRRP entry, click Create. You see the Create VRRP IP Addresses window. |
| Step 4 | Complete the fields in this window to create a new VRRP IP address, and click OK or Apply. |
-

Setting the Priority for the Virtual Router

The valid range to assign a virtual router priority is 1 to 254 with 1 being the lowest priority and 254 being the highest priority. The default value is 100 for switches with secondary IP addresses and 255 for switches with the primary IP address.

Setting the Time Interval for Advertisement Packets

The valid time range for an advertisement packet on an interface using IPv4 is between 1 and 255 seconds. The default value is 1 (one) second. If the switch has the primary IP address, this time must be specified.

Configuring or Enabling Priority Preemption

You can enable a higher-priority backup virtual router to preempt the lower-priority master virtual router.



Note	If the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.
-------------	------------------------------------------------------------------------------------------------------------



Note	The VRRP preemption is not supported on IP storage Gigabit Ethernet interfaces.
-------------	---------------------------------------------------------------------------------

Setting Virtual Router Authentication

VRRP security provides three options, including simple text authentication, MD5 authentication, and no authentication.

- Simple text authentication uses a unique, 1 to 8 character password that is used by all switches participating in the same virtual router. This password should be different from other security passwords.
- MD5 authentication uses a unique, 16 character key that is shared by all switches participating in the same virtual router. This secret key is shared by all switches in the same virtual router.
- No authentication is the default option.

You can configure the key using the authentication option in the VRRP submode and distribute it using the configuration file. The security parameter index (SPI) settings assigned in this option should be unique for each VSAN.



Note All VRRP configurations must be duplicated.



Note VRRP router authentication does not apply to IPv6.

Tracking the Interface Priority

Interface state tracking changes the priority of the virtual router based on the state of another interface in the switch. When the tracked interface is down, the priority reverts to the priority value for the virtual router (see the [Setting the Priority for the Virtual Router, on page 893](#)). When the tracked interface is up, the priority of the virtual router is restored to the interface state tracking value. You can track the state of either a specified VSAN interface or the management interface (mgmt 0). The interface state tracking feature is disabled by default.



Note For interface state tracking to function, you must enable preemption on the interface.

Field Descriptions for IP Services

This section describes the field descriptions.

IP Routes

Field	Description
Routing Enabled	When this check box is enabled, the switch is acting as in IP router.
Destination, Mask, Gateway	The value that identifies the local interface through which the next hop of this route should be reached.
Metric	The primary routing metric for this route.

Field	Description
Interface	The local interface through which the next hop of this route should be reached.
Active	Indicates whether the route is active.

IP Statistics ICMP

Field	Description
InParmProbs	The number of ICMP Parameter Problem messages received.
OutParmProbs	The number of ICMP Parameter Problem messages sent.
InSrcQuenchs	The number of ICMP Source Quench messages received.
InRedirects	The number of ICMP Redirect messages received.
InEchos	The number of ICMP Echo (request) messages received.
InEchoReps	The number of ICMP Echo Reply messages received.
InTimestamps	The number of ICMP Timestamp (request) messages received.
InTimestampReps	The number of ICMP Timestamp Reply messages received.
InAddrMasks	The number of ICMP Address Mask Request messages received.
InAddrMaskReps	The number of ICMP Address Mask Reply messages received.
InDestUnreachs	The number of ICMP Destination Unreachable messages received.
InTimeExcds	The number of ICMP Time Exceeded messages received.
OutSrcQuenchs	The number of ICMP Source Quench messages sent.
OutRedirects	The number of ICMP Redirect messages sent. For a host, this value will always be N/A, since hosts do not send redirects.
OutEchos	The number of ICMP Echo (request) messages sent.
OutEchoReps	The number of ICMP Echo Reply messages sent.
OutTimestamps	The number of ICMP Timestamp (request) messages sent.
OutTimestampReps	The number of ICMP Timestamp Reply messages sent.
OutAddrMasks	The number of ICMP Address Mask Request messages sent.
OutAddrMaskReps	The number of ICMP Address Mask Reply messages sent.
OutDestUnreachs	The number of ICMP Destination Unreachable messages sent.
OutTimeExcds	The number of ICMP Time Exceeded messages sent.

IP Statistics IP

Field	Description
InHdrErrors	The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.
InAddrErrors	The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. For entities that are not IP routers, and do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
InUnknownProtos	The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
InDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (for example, for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.
OutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (for example, for lack of buffer space). Note that this counter would include datagrams counted in ipForwDatagrams if any such frames met this (discretionary) discard criterion.
OutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any frames counted in ipForwDatagrams which meet this no-route criterion. Note that this includes any datagrams which a host cannot route because all of its default routers are down.
FragFails	The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be, for example, because their Don't Fragment flag was set.
ReasmFails	The number of failures detected by the IP reassembly algorithm (for example, timed out, errors, etc). Note that this is not necessarily a count of discarded IP fragments since some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received.
InReceives	The total number of input datagrams received from interfaces, including those received in error.
InDelivers	The total number of input datagrams successfully delivered to IP user protocols (including ICMP).
OutRequests	The total number of IP datagrams which local IP user protocols (including ICMP) supplied to IP in requests for transmission. Note that this counter does not include any datagrams counted in ipForwDatagrams.

Field	Description
ForwDatagrams	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP routers, this counter will include only those frames which were source-routed via this entity, and the Source-Route option processing was successful.
FragOKs	The number of IP datagrams that have been successfully fragmented at this entity.
FragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
ReasmReqds	The number of IP fragments received which needed to be reassembled at this entity.
ReasmOKs	The number of IP datagrams successfully reassembled.

IP Statistics SNMP

Field	Description
BadVersions	The total number of SNMP messages which were delivered to the SNMP entity and were for an unsupported SNMP version.
BadCommunityNames	The total number of SNMP messages delivered to the SNMP entity which used a SNMP community name not known to said entity.
BadCommunityUses	The total number of SNMP messages delivered to the SNMP entity which represented an SNMP operation which was not allowed by the SNMP community named in the message.
ASNParseErrs	The total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages.
TooBig	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is tooBig.
SilentDrops	The total number of GetRequest-PDUs, GetNextRequest-PDUs, GetBulkRequest-PDUs, SetRequest-PDUs, and InformRequest-PDUs delivered to the SNMP entity which were silently dropped because the size of a reply containing an alternate Response-PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the request.
ProxyDrops	The total number of GetRequest-PDUs, GetNextRequest-PDUs, GetBulkRequest-PDUs, SetRequest-PDUs, and InformRequest-PDUs delivered to the SNMP entity which were silently dropped because the transmission of the (possibly translated) message to a proxy target failed in a manner (other than a timeout) such that no Response-PDU could be returned.
NoSuchNames	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is noSuchName.

Field	Description
BadValues	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is badValue.
ReadOnlys	The total number valid SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is readOnly. It should be noted that it is a protocol error to generate an SNMP PDU which contains the value readOnly in the error-status field, as such this is provided as a means of detecting incorrect implementations of the SNMP.
GenErrs	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is genErr.
Pkts	The total number of messages delivered to the SNMP entity from the transport service.
GetRequests	The total number of SNMP Get-Request PDUs which have been accepted and processed by the SNMP protocol entity.
GetNexts	The total number of SNMP Get-Next PDUs which have been accepted and processed by the SNMP protocol entity.
SetRequests	The total number of SNMP Set-Request PDUs which have been accepted and processed by the SNMP protocol entity.
OutTraps	The total number of SNMP Trap PDUs which have been generated by the SNMP protocol entity.
OutGetResponses	The total number of SNMP Get-Response PDUs which have been generated by the SNMP protocol entity.
OutPkts	The total number of SNMP Messages which were passed from the SNMP protocol entity to the transport service.
TotalReqVars	The total number of MIB objects which have been retrieved successfully by the SNMP protocol entity as the result of receiving valid SNMP Get-Request and Get-Next PDUs.
TotalSetVars	The total number of MIB objects which have been altered successfully by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs.

IP Statistics UDP

Field	Description
InErrors	The number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port.
InDatagrams	The total number of UDP datagrams delivered to UDP users.
OutDatagrams	The total number of UDP datagrams sent from this entity.

Field	Description
NoPorts	The total number of received UDP datagrams for which there was no application at the destination port.

mgmt0 Statistics

Field	Description
InErrors	Total number of received errors on the interface.
OutErrors	Total number of transmitted errors on the interface.
InDiscards	Total number of received discards on the interface.
OutDiscards	Total number of transmitted discards on the interface.
RxBytes	Total number of bytes received.
TxBytes	Total number of bytes transmitted.
RxFrames	Total number of frames received.
TxFrames	Total number of frames transmitted.

TCP UDP TCP

Field	Description
State	The state of this TCP connection.

TCP UDP UDP

Field	Description
Port	The local port number for this UDP listener.

VRRP General

Field	Description
IP Address Type, Vrid, Interface	The IP address type (IPv4, IPv6, or DNS), the virtual router ID, and the interface.
Admin	The admin state of the virtual router (active or notInService).

Field	Description
Oper	The current state of the virtual router. There are three defined values: <ul style="list-style-type: none"> • initialize— Indicates that all the virtual router is waiting for a startup event. • backup— Indicates the virtual router is monitoring the availability of the master router. • master— Indicates that the virtual router is forwarding frames for IP addresses that are associated with this router.
Priority	Specifies the priority to be used for the virtual router master election process. Higher values imply higher priority. A priority of 0 is sent by the master router to indicate that this router has ceased to participate in VRRP and a backup virtual router should transition to become a new master. A priority of 255 is used for the router that owns the associated IP address(es).
AdvInterval	The time interval, in seconds, between sending advertisement messages. Only the master router sends VRRP advertisements.
PreemptMode	Controls whether a higher priority virtual router will preempt a lower priority master.
UpTime	When this virtual router transitioned out of initialized.
Version	The VRRP version on which this VRRP instance is running.
AcceptMode	Controls whether a virtual router in Master state will accept packets addressed to the address owner's IPv6 address as its own if it is not the IPv6 address owner. If true, the virtual router in Master state will accept. If false, the virtual router in Master state will not accept.

VRRP IP Addresses

Field	Description
Interface, VRRP ID, IP Address	Interface, Virtual Router Redundancy Protocol ID, and associated IP address.

VRRP Statistics

Field	Description
IP Address Type, VrId, Interface	The IP address type (IPv4, IPv6, or DNS), the virtual router ID, and the interface.
LastAdvRx	The total number of VRRP advertisements received by this virtual router.
Protocol Traffic MasterIpAddr	The master router's real (primary) IP address. This is the IP address listed as the source in VRRP advertisement last received by this virtual router.
Protocol Traffic BecomeMaster	The total number of times that this virtual router's state has transitioned to MASTER.

Field	Description
Priority 0 Rx	The total number of VRRP frames received by the virtual router with a priority of 0.
Priority 0Tx	The total number of VRRP frames sent by the virtual router with a priority of 0.
AuthErrors InvalidType	The total number of frames received with an unknown authentication type.
Other Errors dvIntervalErrors	The total number of VRRP advertisement frames received for which the advertisement interval is different than the one configured for the local virtual router.
Other Errors IpTtlErrors	The total number of VRRP frames received by the virtual router with IP TTL (time-to-live) not equal to 255.
Other Errors InvalidTypePktsRcvd	The number of VRRP frames received by the virtual router with an invalid value in the type field.
Other Errors AddressListErrors	The total number of frames received for which the address list does not match the locally configured list for the virtual router.
OtherErrors PacketLengthErrs	The total number of frames received with a frame length less than the length of the VRRP header.
RefreshRate	The interval of time between refreshes.

CDP General

Field	Description
Enable	Whether the Cisco Discovery Protocol is currently running. Entries in CacheTable are deleted when CDP is disabled.
MessageInterval	The interval at which CDP messages are to be generated. The default value is 60 seconds.
HoldTime	The time for the receiving device holds CDP message. The default value is 180 seconds.
LastChange	When the cache table was last changed.

CDP Neighbors

Field	Description
Switch	The Internet address for this entity.
Local Interface	A unique value that identifies the interface on this FCIP device to which this link pertains.
DeviceName	The remote device's name. By convention, it is the device's fully qualified domain name.

Field	Description
DeviceID	The device ID string as reported in the most recent CDP message.
DevicePlatform	The version string as reported in the most recent CDP message.
Interface	The port ID string as reported in the most recent CDP message.
IPAddress	The (first) network-layer address of the device's SNMP-agent as reported in the address TLV of the most recently received CDP message.
NativeVLAN	The remote device's interface's native VLAN, as reported in the most recent CDP message. The value 0 indicates no native VLAN field (TLV) was reported in the most recent CDP message.
PrimaryMgmtAddr	Indicates the (first) network layer address at which the device will accept SNMP messages as reported in the most recently received CDP message.
SecondaryMgmtAddr	Indicates the alternate network layer address at which the device will accept SNMP messages as reported in the most recently received CDP message.

iSNS Profiles

Field	Description
Addr	The address of the iSNS server.
Port	The TCP port of the iSNS server.

iSNS Servers

Field	Description
Name	The name of the iSNS server.
TcpPort	The TCP port used for iSNS messages. If TCP is not supported by this server, the value is 0.
Uptime	The time the server has been active.
ESI Non Response Threshold	The number of ESI messages that will be sent without receiving a response before an entity is unregistered from the iSNS database.
# Entities	The number of entities registered in iSNS on the server.
# Portals	The number of portals registered in iSNS on the server.
# Portal Groups	The number of portal groups registered in iSNS on the server.
# iSCSI Devices	The number of iSCSI Nodes registered in iSNS on the server.

iSNS Entities

Field	Description
Entity ID	The iSNS entity identifier for the entity.
Last Accessed	The time the entity was last accessed.

iSNS Cloud Discovery

Field	Description
AutoDiscovery	Whether automatic cloud discovery is turned on or off.
DiscoveryDelay	Time duration between successive IP cloud discovery runs.
Discovery	<p>The IP network discovery command to be executed.</p> <ul style="list-style-type: none"> all- Run IP network discovery for all the gigabit Ethernet interfaces in the fabric. noOp (default)- No operation is performed.
CommandStatus	<p>The status of the license install / uninstall / update operation.</p> <ul style="list-style-type: none"> success— Discovery operation completed successfully. nProgress— Discovery operation is in progress. none— No discovery operation is performed. NoIpNetworkNameSpecified— IP Cloud name not specified. invalidNetworkName— IP Cloud is not configured. NoIPSPortNameSpecified— Gigabit Ethernet port if index not specified. invalidIPSPortName— Invalid Gigabit Ethernet port interface. generalISNSFailure— General iSNS server failure.

iSNS Clouds

Field	Description
Id	The ID of the IP cloud.
Switch WWN	The WWN of the switch in this table.

iSNS Cloud Interfaces

Field	Description
Name, Switch WWN, Interface, Address	The name, switch WWN, interface, and address of the cloud.

Monitor Dialog Controls

Field	Description
Line Chart	Opens a new window with a line chart representation of the data.
Area Chart	Opens a new window with an area chart representation of the data.
Bar Chart	Opens a new window with a bar chart representation of the data.
Pie Chart	Opens a new window with a pie chart representation of the data.
Reset Cumulative Counters	Resets the counters to 0 if the Column Data display mode is set to Cumulative.
Export to File	Opens a standard Save dialog box. The data is saved as a .TXT file.
Print	Opens a standard Print dialog box.
Update Frequency	The interval at which the data is updated in the monitor dialog.
Column Data	<p>Specifies the type of data that is displayed in the monitor dialog.</p> <ul style="list-style-type: none"> • Absolute Value— Displays the total amount since the switch was booted. This is the default for error monitoring. • Cumulative—Displays the total amount since the dialog was opened. You can reset the counters by clicking the Reset Cumulative Counters button, to gather a new set of cumulative data. • Minimum/sec— Displays the minimum value per second at every refresh interval. • Maximum/sec— Displays the maximum value per second at every refresh interval. • Last Value/sec— Displays the most recent value per second at every refresh interval. This is the default setting for traffic monitoring.
Elapsed	The amount of time that has elapsed since the dialog was opened. You can reset this counter by clicking the Reset Cumulative Counters button, to gather a new set of cumulative data.

iSNS Details iSCSI Nodes

Field	Description
Name	The iSCSI name of the initiator or target associated with the storage node.
Type	The Node Type bit-map defining the functions of this iSCSI node, where 31 is a Target, 30 is an Initiator, 29 is a Control, and all others are reserved.
Alias	The Alias name of the iSCSI node.
ScnBitmap	The State Change Notification (SCN) bitmap for a node.
WWN Token	An optional globally unique 64-bit integer value that can be used to represent the world wide node name of the iSCSI device in a Fibre Channel fabric.

Field	Description
AuthMethod	The iSCSI authentication method enabled for this iSCSI node.

iSNS Details Portals

Field	Description
Addr	The Internet address for this portal.
TcpPort	The port number for this portal.
SymName	The optional Symbolic Name for this portal.
EsiInterval	The Entity Status Inquiry (ESI) Interval for this portal.
TCP ESI	The TCP port number used for ESI monitoring.
TCP Scn	The TCP port used to receive SCN messages from the iSNS server.
SecurityInfo	Security attribute settings for the portal as registered in the Portal Security Bitmap attribute.



CHAPTER 42

Configuring IP Storage

- [Configuring IP Storage, on page 907](#)

Configuring IP Storage

Cisco MDS 9000 Family IP storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch connects separated SAN islands using Fibre Channel over IP (FCIP), and it allows IP hosts to access Fibre Channel storage using the iSCSI protocol.



Note FCIP and iSCSI features are specific to the IPS module and are available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors. The Cisco MDS 9216I switch and the 14/2 Multiprotocol Services (MPS-14/2) module also allow you to use Fibre Channel, FCIP, and iSCSI features. The MPS-14/2 module is available for use in any switch in the Cisco MDS 9200 Series or Cisco MDS 9500 Series.

This chapter includes the following topics:

Information About IP Storage

The IP Storage services module (IPS module) and the MPS-14/2 module allow you to use FCIP and iSCSI features. FCIP and iSCSI features are specific to the IPS module and are available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors. The switch connects separated SAN islands using Fibre Channel over IP (FCIP), and it allows IP hosts to access Fibre Channel storage using the iSCSI protocol.

- **FCIP**—FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices.
- **iSCSI**—The IPS module provides IP hosts access to Fibre Channel storage devices. The IP host sends SCSI commands encapsulated in iSCSI protocol data units (PDUs) to a Cisco MDS 9000 Family switch IPS port over a TCP/IP connection. At this point, the commands are routed from an IP network into a Fibre Channel network and forwarded to the intended target.

The IP Storage services module (IPS module) and the MPS-14/2 module allow you to use FCIP and iSCSI features. Both modules integrate seamlessly into the Cisco MDS 9000 Family, and support the full range of features available on other switching modules, including VSANs, security, and traffic management. The following types of storage services modules are currently available for use in any switch in the Cisco MDS 9200 Series or in the Cisco MDS 9500 Series:

- The 4-port, hot-swappable IPS module (IPS-4) has four Gigabit Ethernet ports.
- The 8-port, hot-swappable IPS module (IPS-8) has eight Gigabit Ethernet ports.
- The MPS-14/2 module has 14 Fibre Channel ports (numbered 1 through 14) and two Gigabit Ethernet ports (numbered 1 and 2).

Gigabit Ethernet ports in these modules can be configured to support the FCIP protocol, the iSCSI protocol, or both protocols simultaneously:

- FCIP—FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices.

[Figure 149: FCIP Scenarios, on page 908](#) shows how the IPS module is used in different FCIP scenarios.

Figure 149: FCIP Scenarios

- iSCSI—The IPS module provides IP hosts access to Fibre Channel storage devices. The IP host sends SCSI commands encapsulated in iSCSI protocol data units (PDUs) to a Cisco MDS 9000 Family switch IPS port over a TCP/IP connection. At this point, the commands are routed from an IP network into a Fibre Channel network and forwarded to the intended target.

[Figure 150: iSCSI Scenarios, on page 908](#) depicts the iSCSI scenarios in which the IPS module is used.

Figure 150: iSCSI Scenarios

This section contains the following topics:

IPS Module Upgrade



Caution

A software upgrade is only disruptive for the IPS module. The NX-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

IPS modules use a rolling upgrade install mechanism where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each IPS module in a switch requires a 5-minute delay before the next IPS module is upgraded.

MPS-14/2 Module Upgrade



Caution

A software upgrade is only partially disruptive for the MPS-14/2 module. The NX-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

The MPS-14/2 modules have 14 Fibre Channel ports (nondisruptive upgrade) and two Gigabit Ethernet ports (disruptive upgrade). MPS-14/2 modules use a rolling upgrade install mechanism for the two Gigabit Ethernet ports where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each MPS-14/2 module in a switch requires a 5-minute delay before the next module is upgraded.

Supported Hardware

You can configure the FCIP and iSCSI features using one or more of the following hardware:

- IPS-4 and IPS-8 modules (refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for more information)
- MPS-14/2 module (refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for more information).



Note In both the MPS-14/2 module and the Cisco MDS 9216i integrated supervisor module, the port numbering differs for the Fibre Channel ports and the Gigabit Ethernet ports. The Fibre Channel ports are numbered from 1 through 14 and the Gigabit Ethernet ports are numbered 1 and 2.

- Cisco MDS 9216i Switch (refer to the *Cisco MDS 9200 Series Hardware Installation Guide*).

Gigabit Ethernet Interfaces for IPv4 Configuration

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each IPS module or MPS-14/2 module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured.

A new port mode, called IPS, is defined for Gigabit Ethernet ports on each IPS module or MPS-14/2 module. IP storage ports are implicitly set to IPS mode, so it can be used to perform only iSCSI and FCIP storage functions. IP storage ports do not bridge Ethernet frames or route other IP packets.

Each IPS port represents a single virtual Fibre Channel host in the Fibre Channel SAN. All iSCSI hosts connected to this IPS port are merged and multiplexed through the single Fibre Channel host.



Note For information about configuring FCIP, see the *Configuring FCIP* chapter. For information about configuring iSCSI, see the *Configuring iSCSI* chapter.

In large scale iSCSI deployments where the Fibre Channel storage subsystems require explicit LUN access control for every host device, use of proxy-initiator mode simplifies the configuration.



Note The Gigabit Ethernet interfaces on the MPS-14/2 module do not support EtherChannel.




Note To configure IPv6 on a Gigabit Ethernet interface, see the Security Configuration Guide, Cisco DCNM for SAN.



Tip Gigabit Ethernet ports on any IPS module or MPS-14/2 module should not be configured in the same Ethernet broadcast domain as the management Ethernet port—they should be configured in a different broadcast domain, either by using separate standalone hubs or switches or by using separate VLANs.

Basic Gigabit Ethernet Configuration

The following figure shows an example of a basic Gigabit Ethernet IP version 4 (IPv4) configuration.

Figure 151: Gigabit Ethernet IPv4 Configuration Example

The port on the Ethernet switch to which the Gigabit Ethernet interface is connected should be configured as a host port (also known as access port) instead of a switch port. Spanning tree configuration for that port (on the Ethernet switch) should be disabled. This helps avoid the delay in the management port coming up due to delay from Ethernet spanning tree processing that the Ethernet switch would run if enabled. For Cisco Ethernet switches, use either the **switchport host** command in Cisco IOS or the **set port host command** in Catalyst OS.

IPS Module Core Dumps

IPS core dumps are different from the system's kernel core dumps for other modules. When the IPS module's operating system (OS) unexpectedly resets, it is useful to obtain a copy of the memory image (called a IPS core dump) to identify the cause of the reset. Under that condition, the IPS module sends the core dump to the supervisor module for storage. Cisco MDS switches have two levels of IPS core dumps:

- Partial core dumps (default)—Each partial core dump consists of four parts (four files). All four files are saved in the active supervisor module.

Use the **show cores** command to list these files.

- Full core dumps—Each full core dump consists of 75 parts (75 files). The IPS core dumps for the MPS-14/2 module and the Cisco MDS 9216i Switch only contains 38 parts. This dump cannot be saved on the supervisor module because of its large space requirement. They are copied directly to an external TFTP server.

Use the **system cores tftp:** command to configure an external TFTP server to copy the IPS core dump (and other core dumps).

Interface Descriptions Configuration

See the Interfaces Configuration Guide, Cisco DCNM for SAN for details on configuring the switch port description for any interface.

Beacon Mode Configuration

See the Interfaces Configuration Guide, Cisco DCNM for SAN for details on configuring the beacon mode for any interface.

Autonegotiation Configuration

By default, autonegotiation is enabled all Gigabit Ethernet interface. You can enable or disable autonegotiation for a specified Gigabit Ethernet interface. When autonegotiation is enabled, the port automatically detects the speed or pause method, and duplex of incoming signals based on the link partner. You can also detect link up conditions using the autonegotiation feature.

MTU Frame Size Configuration

You can configure the interfaces on a switch to transfer large (or jumbo) frames on a port. The default IP maximum transmission unit (MTU) frame size is 1500 bytes for all Ethernet ports. By configuring jumbo frames on a port, the MTU size can be increased up to 9000 bytes.



Note The minimum MTU size is 576 bytes.



Tip MTU changes are disruptive, all FCIP links and iSCSI sessions flap when the software detects a change in the MTU size.

Promiscuous Mode Configuration

You can enable or disable promiscuous mode on a specific Gigabit Ethernet interface. By enabling the promiscuous mode, the Gigabit Ethernet interface receives all the packets and the software then filters and discards the packets that are not destined for that Gigabit Ethernet interface.

About VLANs for Gigabit Ethernet

Virtual LANs (VLANs) create multiple virtual Layer 2 networks over a physical LAN network. VLANs provide traffic isolation, security, and broadcast control.

Gigabit Ethernet ports automatically recognize Ethernet frames with IEEE 802.1Q VLAN encapsulation. If you need to have traffic from multiple VLANs terminated on one Gigabit Ethernet port, configure subinterfaces—one for each VLAN.

If the IPS module or MPS-14/2 module is connected to a Cisco Ethernet switch, and you need to have traffic from multiple VLANs coming to one IPS port, verify the following requirements on the Ethernet switch:

- The Ethernet switch port connected to the IPS module or MPS-14/2 module is configured as a trunking port.
- The encapsulation is set to 802.1Q and not ISL, which is the default.

Use the VLAN ID as a subscription to the Gigabit Ethernet interface name to create the subinterface name: slot-number / port-numberVLAN-ID .

Interface Subnet Requirements

Gigabit Ethernet interfaces (major), subinterfaces (VLAN ID), and management interfaces (mgmt 0) can be configured in the same or different subnet depending on the configuration (see [Table 110: Subnet Requirements for Interfaces](#) , on page 911).

Table 110: Subnet Requirements for Interfaces

Interface 1	Interface 2	Same Subnet Allowed	Notes
Gigabit Ethernet 1/1	Gigabit Ethernet 1/2	Yes	Two major interfaces can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.100	Yes	Two subinterfaces with the same VLAN ID can be configured in the same or different subnets.

Interface 1	Interface 2	Same Subnet Allowed	Notes
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.200	No	Two subinterfaces with different VLAN IDs cannot be configured in the same subnet.
Gigabit Ethernet 1/1	Gigabit Ethernet 1/1.100	No	A subinterface cannot be configured on the same subnet as the major interface.
mgmt0	Gigabit Ethernet 1/1.100	No	The mgmt0 interface cannot be configured in the same subnet as the Gigabit Ethernet interfaces or subinterfaces.
mgmt0	Gigabit Ethernet 1/1	No	

**Note**

The configuration requirements in [Table 110: Subnet Requirements for Interfaces](#), on page 911 also apply to Ethernet PortChannels.

Verifying Gigabit Ethernet Connectivity

Once the Gigabit Ethernet interfaces are connected with valid IP addresses, verify the interface connectivity on each switch. Ping the IP host using the IP address of the host to verify that the static IP route is configured correctly.

**Note**

If the connection fails, verify the following, and ping the IP host again:- The IP address for the destination (IP host) is correctly configured.- The host is active (powered on). - The IP route is configured correctly. - The IP host has a route to get to the Gigabit Ethernet interface subnet. - The Gigabit Ethernet interface is in the up state.

Gigabit Ethernet High Availability

Virtual Router Redundancy Protocol (VRRP) and Ethernet PortChannels are two Gigabit Ethernet features that provide high availability for iSCSI and FCIP services.

VRRP for iSCSI and FCIP Services

VRRP provides a redundant alternate path to the Gigabit Ethernet port for iSCSI and FCIP services. VRRP provides IP address failover protection to an alternate Gigabit Ethernet interface so the IP address is always available (see [Figure 152: VRRP Scenario](#), on page 912).

Figure 152: VRRP Scenario

All members of the VRRP group (see [Figure 152: VRRP Scenario](#), on page 912) must be IP storage Gigabit Ethernet ports. VRRP group members can be one or more of the following interfaces:

- One or more interfaces in the same IPS module or MPS-14/2 module
- Interfaces across IPS modules or MPS-14/2 modules in one switch

- Interfaces across IPS modules or MPS-14/2 modules in different switches
- Gigabit Ethernet subinterfaces
- Ethernet PortChannels and PortChannel subinterfaces



Note You can configure no more than seven VRRP groups, both IPv4 and IPv6, on a Gigabit Ethernet interface, including the main interface and all subinterfaces.

About Ethernet PortChannel Aggregation

Ethernet PortChannels refer to the aggregation of multiple physical Gigabit Ethernet interfaces into one logical Ethernet interface to provide link redundancy and, in some cases, higher aggregated bandwidth and load balancing.

An Ethernet switch connecting to the MDS switch Gigabit Ethernet port can implement load balancing based on the IP address, IP address and UDP/TCP port number, or MAC address. Due to the load balancing scheme, the data traffic from one TCP connection is always sent out on the same physical Gigabit Ethernet port of an Ethernet PortChannel. For the traffic coming to the MDS, an Ethernet switch can implement load balancing based on its IP address, its source-destination MAC address, or its IP address and port. The data traffic from one TCP connection always travels on the same physical links. To make use of both ports for the outgoing direction, multiple TCP connections are required.

All FCIP data traffic for one FCIP link is carried on one TCP connection. Consequently, the aggregated bandwidth is 1 Gbps for that FCIP link.



Note The Cisco Ethernet switch's PortChannel should be configured as a static PortChannel, and not the default 802.3ad protocol.

Ethernet PortChannels can only aggregate two physical interfaces that are adjacent to each other on a given IPS module (see [Figure 153: Ethernet PortChannel Scenario, on page 913](#)).



Note PortChannel members must be one of these combinations: ports 1–2, ports 3–4, ports 5–6, or ports 7–8.

Figure 153: Ethernet PortChannel Scenario

In [Figure 153: Ethernet PortChannel Scenario, on page 913](#), Gigabit Ethernet ports 3 and 4 in slot 9 are aggregated into an Ethernet PortChannel. Ethernet PortChannels are not supported on MPS-14/2 modules and 9216i IPS modules.



Note PortChannel interfaces provide configuration options for both Gigabit Ethernet and Fibre Channel. However, based on the PortChannel membership, only Gigabit Ethernet parameters or Fibre Channel parameters are applicable.

CDP

The Cisco Discovery Protocol (CDP) is an advertisement protocol used by Cisco devices to advertise itself to other Cisco devices in the same network. CDP runs on the data link layer and is independent of Layer 3 protocols. CDP is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interfaces on the IPS and MPS-14/2 modules.

CDP version 1 (v1) and version 2 (v2) are supported in Cisco MDS 9000 Family switches. CDP packets with any other version number are silently discarded when received.

See the Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide.

Licensing Requirements for IP Storage

The following table shows the licensing requirements for this feature:

License	License Description
SAN extension over IP package for IPS-8 modules • (SAN_EXTN_OVER_IP) SAN extension over IP package for IPS-4 modules • (SAN_EXTN_OVER_IP_IPS4)	The following features apply to IPS-8 and IPS-4 modules: <ul style="list-style-type: none"> • FCIP • FCIP compression • FCIP write acceleration • FCIP tape read acceleration • SAN extension tuner features • IVR over FCIP • IVR NAT over FCIP • Network Stimulator
SAN extension over IP package for MPS-14/2 modules • (SAN_EXTN_OVER_IP_IPS2)	The following features apply to the MPS-14/2 module and the fixed Cisco MDS 9216i Switch IP ports: <ul style="list-style-type: none"> • FCIP • Hardware-based FCIP compression • FCIP write acceleration • FCIP tape read acceleration • SAN extension tuner features • IVR over FCIP • IVR NAT over FCIP

Guidelines and Limitations



Tip

If IPv4-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group.

Follow these guidelines when configuring IPv4-ACLs for Gigabit Ethernet interfaces:

- Only use Transmission Control Protocol (TCP) or Internet Control Message Protocol (ICMP).



Note Other protocols such as User Datagram Protocol (UDP) and HTTP are not supported in Gigabit Ethernet interfaces. Applying an ACL that contains rules for these protocols to a Gigabit Ethernet interface is allowed but those rules have no effect.

- Apply IPv4-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
 - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
 - The **established**, **precedence**, and **fragments** options are ignored when you apply IPv4-ACLs (containing these options) to Gigabit Ethernet interfaces.
 - If an IPv4-ACL rule applies to a preexisting TCP connection, that rule is ignored. For example if there is an existing TCP connection between A and B, and an IPv4-ACL specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.

Default Settings

Table 111: Default Gigabit Ethernet Parameters , on page 915 lists the default settings for IP storage services parameters.

Table 111: Default Gigabit Ethernet Parameters

Parameters	Default
IPS core size	Partial

Configuring IP Storage

This section includes the following topics:

Configuring IPS Core Dumps

To configure the Gigabit Ethernet interface, follow these steps:

Procedure

- Step 1** From Cisco DCNM-SAN, choose Switches > Interfaces > Gigabit Ethernet in the Physical Attributes pane. You see the Gigabit Ethernet configuration in the Information pane.

From Device Manager, right-click the Gigabit Ethernet port that you want to configure and choose Configure.... You see the Gigabit Ethernet configuration dialog box.
- Step 2** Click the General tab in Cisco DCNM-SAN, or click the GigE tab in Device Manager to display the general configuration options for the interface.
- Step 3** Set the description and MTU value for the interface. The valid value for the MTU field can be a number in the range from 576 to 9000.

- Step 4** Set Admin up or down and check the CDP check box if you want this interface to participate in CDP.
- Step 5** Set IpAddress/Mask with the IP address and subnet mask for this interface.
- Step 6** From Cisco DCNM-SAN, click the Apply Changes icon to save these changes, or click the Undo Changes icon to discard changes.
- From Device Manager, click Apply to save these changes, or click Close to discard changes and close the Gigabit Ethernet configuration dialog box.

Verifying IP Storage Configuration

To display IP storage configuration information, perform one of the following tasks:

Command	Purpose
show module	Verifies the status of the module.
show interface gigabitethernet 8/1	Displays the gigabit ethernet interface.
show interface gigabitethernet 4/2.100	Displays the gigabit ethernet subinterface.
show ips stats mac interface gigabitethernet 8/1	Displays ethernet MAC statistics.
show ips stats dma-bridge interface gigabitethernet 7/1	Displays DMA-Bridge statistics.
show ips stats tcp interface gigabitethernet 4/1	Displays TCP statistics.
show ips stats tcp interface gigabitethernet 4/1 detail	Displays Detailed TCP statistics.
show ips stats icmp interface gigabitethernet 2/1	Displays ICMP statistics.

This section includes the following topics:

Verifying Module Status

To verify the status of the module, follow these steps:

Procedure

- Step 1** Select a switch in the Fabric pane.
- Step 2** Open the **Switches** folder and select **Hardware** in the Physical Attributes pane.
- You see the status for all modules in the switch in the Information pane.

Field Descriptions for IP Storage

This section describes the following field descriptions.

FCIP Profiles

Field	Description
IP Address	The Internet address for this entity.
Port	A TCP port other than the FCIP well-known port on which the FCIP entity listens for new TCP connection requests.
SACK	Whether the TCP Selective Acknowledgement Option is enabled to allow the receiver end to acknowledge multiple lost frames in a single ACK, enabling faster recovery.
KeepAlive (s)	The TCP keepalive timeout for all links within this entity.
ReTrans MinTimeout (ms)	The TCP minimum retransmit timeout for all the links on this entity.
ReTrans Max	The maximum number of times that the same item of data will be retransmitted over a TCP connection. If delivery is not acknowledged after this number of retransmissions then the connection is terminated.
Send BufSize (KB)	The aggregate TCP send window for all TCP connections on all links within this entity. This value is used for egress flow control. When the aggregate of the data queued on all connections within this entity reaches this value, the sender is flow controlled.
Bandwidth Max (Kb)	This is an estimate of the bandwidth of the network pipe used for the B-D product computation, which lets us derive the TCP receive window to advertise.
Bandwidth Min (Kb)	The minimum available bandwidth for the TCP connections on the links within this entity.
Est Round Trip Time (us)	This is an estimate of the round trip delay of the network pipe used for the B-D product computation, which lets us derive the TCP receive window to advertise.
PMTU Enable	The path MTU discovery.
PMTU ResetTimeout (sec)	The time interval for which the discovered path MTU is valid, before MSS reverts back to the negotiated TCP value.
CWM Enable	If true, congestion window monitoring is enabled.
CWM BurstSize (KB)	The maximum burst sent after a TCP sender idle period.
Max Jitter	The maximum delay variation (not due to congestion) that can be experienced by TCP connections on this interface.

FCIP Tunnels

Field	Description
Interface	This identifies the interface on this FCIP device to which this link pertains.
Attached	The interface on which this FCIP link was initiated.

Field	Description
B Port Enable	If true, the B port mode is enabled on the local FCIP link.
B Port KeepAlive	If true, a message is sent in response to a (Fibre Channel) ELS Echo frame received from the peer. Some B Port implementations use ELS Echo request/response frames as Link Keep Alive.
Remote IP Address	The Internet address for the remote FCIP entity.
Remote TCP Port	The remote TCP port to which the local FCIP entity will connect if and when it initiates a TCP connection setup for this link.
Spc Frames Enable	If true, the TCP active opener initiates FCIP special frames and the TCP passive opener responds to the FCIP special frames. If it is set to false, the FCIP special frames are neither generated nor responded to.
Spc Frames RemoteWWN	The world wide name of the remote FC fabric entity. If this is a zero length string then this link would accept connections from any remote entity. If a WWN is specified then this link would accept connections from a remote entity with this WWN.
Spc Frames Remote Profile Id	The remote FCIP entity's identifier.

FCIP Tunnels (Advanced)

Field	Description
Interface	The interface on which this FCIP link was initiated.
Timestamp Enable	If true, the timestamp in FCIP header is to be checked.
Timestamp Tolerance	The accepted time difference between the local time and the timestamp value received in the FCIP header. By default this value will be EDTOV/2. EDTOV is the Error_Detect_Timeout Value used for Fibre Channel ports as the timeout value for detecting an error condition.
Number Connections	The maximum number of TCP connections allowed on this link.
Passive	If false, this link endpoint actively tries to connect to the peer. If true, the link endpoint waits for the peer to connect to it.
QoS Control	The value to be set for the ToS field in IP header for the TCP control connection.
QoS Data	The value to be set for the ToS field in IP header for the TCP data connection.
IP Compression	What algorithm is used, if any.
Write Accelerator	The write accelerator allows for enhancing SCSI write performance.
Tape Accelerator	If true, the tape accelerator (which allows for enhancing Tape write performance) is enabled.
Tape Accelerator Oper	Write acceleration is enabled for the FCIP link.

Field	Description
TapeRead Accelerator Oper	Enabled automatically when the tape accelerator oper is active.
FlowCtrlBufSize Tape (KB)	The size of the flow control buffer (64 K to 32 MB). If set to 0, flow control buffer size is calculated automatically by the switch.
IPSec	Indicates whether the IP security has been turned on or off on this link.
XRC Emulator	Check to enable XRC emulator. It is disabled by default.
XRC Emulator Oper	Indicates the operational status of XRC emulator.

FCIP Tunnels (FICON TA)

Field	Description
Interface	A unique value that identifies the interface on this FCIP device to which this link pertains.
VSAN List Admin	The list of VSANs for which FICON tape acceleration is configured.
VSAN List Oper	The list of VSANs for which FICON tape acceleration is operationally on.

FCIP Tunnels Statistics

Field	Description
Interface	A unique value that identifies the interface on this FCIP device to which this link pertains.
Rx IPCompRatio	The IP compression ratio for received packets on the FCIP device. The value of this object will be presented as a floating point number with two digits after the decimal point.
Tx IPCompRatio	The IP compression ratio for transmitted packets on the FCIP device. The value of this object will be presented as a floating point number with two digits after the decimal point.

FCIP XRC Statistics

Field	Description
ProfileId	Unique ID of the profile.
Interface	Name of the interface.
RRSAccelerated	The number of read record set IUs accelerated.
RRSForwarded	Number of read record set IUs forwarded.
BusyStatus	Number of instances of busy status received from the control unit.
UnitCheckStatus	Number of instances of unit check status received from the control unit.
cfmFeipLinkExtXRCEStatsSelReset	Number of selective resets processed.

Field	Description
BufferAllocErrors	Number of buffer allocation errors.

iSCSI Connection

Field	Description
LocalAddr	The local Internet network address used by this connection.
RemoteAddr	The remote Internet network address used by this connection.
CID	The iSCSI connection ID for this connection.
State	The current state of this connection, from an iSCSI negotiation point of view. <ul style="list-style-type: none"> • login— The transport protocol connection has been established, but a valid iSCSI login response with the final bit set has not been sent or received. • full— A valid iSCSI login response with the final bit set has been sent or received. • logout— A valid iSCSI logout command has been sent or received, but the transport protocol connection has not yet been closed.
MaxRecvDSLen	The maximum data payload size supported for command or data PDUs in use within this connection. Note that the size of reported in bytes even though the negotiation is in 512 K blocks.
SendMarker	Indicates whether or not this connection is inserting markers in its outgoing data stream.
HeaderDigest	The iSCSI header digest scheme in use within this connection.
DataDigest	The iSCSI data digest scheme in use within this connection.

iSCSI Initiators

Field	Description
Name or IP Address	A character string that is a globally unique identifier for the node represented by this entry.
VSAN Membership	The list of configured VSANs the node represented by this entry can access.
Dynamic	If true, then the node represented by this entry is automatically discovered.
Initiator Type	Indicates whether the node is a host that participates in iSCSI load-balancing.
Persistent Node WWN	If true, then the same FC address is assigned to the node if it were to be represented again in the FC domain with the same node name. Note that the node FC address is either automatically assigned or manually configured.
SystemAssigned Node WWNN	If true, the FC address is automatically assigned to this node. If false, then the FC address has to be configured manually.
Node WWN	The persistent FC address of the node.

Field	Description
Persistent Port WWN	If true, then the same FC address is assigned to the ports of the node if it were to be represented again in the FC domain with the same node name.
Port WWN	All the FC port addresses associated with this node.
AuthUser	This is the only CHAP user name that the initiator is allowed to log in with.
Target UserName	(Optional) The user name to be used for login. If you do not supply a username, the global user name is used.
Target Password	(Optional) The password to be used for login. If you do not supply a password, the global password is used.
Load Metric	A configured load metric of this iSCSI initiator for the purpose of iSCSI load balancing.
Auto Zone Name	The zone name that is used when the system creates automatic zone for this initiator's specific list of targets.

iSCSI Targets

Field	Description
Dynamically Import FC Targets	Check this option to dynamically import FC targets into the iSCSI domain. A target is not imported if it already exists in the iSCSI domain.
iSCSI Name	The iSCSI name of the node represented by this entry.
Dynamic	Indicates if the node represented by this entry was either automatically discovered or configured manually.
Primary Port WWN	The FC address for this target.
Secondary Port WWN	The optional secondary FC address for this target. This is the FC address used if the primary cannot be reached.
LUN Map iSCSI	The configured default logical unit number of this LU.
LUN Map FC Primary	The logical unit number of the remote LU for the primary port address.
LUN Map FC Secondary	The logical unit number of the remote LU for the secondary port address.
Initiator Access All	If true, then all the initiators can access this target even those which are not in the initiator permit list of this target. If false, then only initiators which are in the permit list are allowed access to this target.
Initiator Access List	Lists all the iSCSI nodes that are permitted to access the node represented by this entry. If AllAllowed is false and the value of List is empty, then no initiators are allowed to access this target.
Advertised Interfaces	Lists all the interfaces on which the target could be advertised.

Field	Description
Trespass Mode	The trespass mode for this node. Every iSCSI target represents one or more port(s) on the FC target. If true, the node instructs the FC node to present all LUN I/O requests to secondary port if the primary port is down.
RevertToPrimaryPort	Indicates if it is required to revert back to primary port if the FC target comes back online.

iSCSI Session Initiators

Field	Description
Name or IP Address	The name or IP address of the initiator port.
Alias	The initiator alias acquired at login.

Module Control

Field	Description
Module Id	ID of the module.
Admin Status	Enables or disables the iSCSI feature for the module.
OperStatus	Shows whether the iSCSI interface is enabled or disabled for the module.

iSCSI Global

Field	Description
AuthMethod	The authentication method.
InitiatorIdleTimeout	The time for which the gateway (representing a FC target) waits from the time of last iSCSI session to a iSCSI initiator went down, before purging the information about that iSCSI initiator.
iSLB ZonesetActivate	Checking this option performs automatic zoning associated with the initiator targets
DynamicInitiator	This field determines how dynamic iSCSI initiators are created. Selecting the iSCSI option (default) creates dynamic iSCSI initiators. If you select iSLB then the an iSLB dynamic initiator is created. Selecting the deny option does not allow dynamic creation of the initiators.
Target UserName	The default user name used for login. If an initiator user name is specified, that user name is used instead.
Target Password	The default password used for login. If an initiator password is specified, that password is used instead.

iSCSI Session Statistics

Field	Description
PDU Command	The count of command PDUs transferred on this session.
PDU Response	The count of response PDUs transferred on this session.
Data Tx	The count of data bytes that were transmitted by the local iSCSI node on this session.
Data Rx	The count of data bytes that were received by the local iSCSI node on this session.
Errors Digest	Authentication errors.
Errors CxnTimeout	Connection timeouts.

iSCSI iSLB VRRP

Field	Description
VrId, IpVersion	The virtual router number and the IP version (IPv4, IPv6, or DNS).
Load Balance	Indicates whether load balancing is enabled.

iSCSI Initiator Access

Field	Description
Initiator Name	The iSCSI node name.

Initiator Specific Target

Field	Description
Name	A globally unique identifier for the node.
Port WWN(s) Primary	The Fibre Channel target's port addresses associated with this iSCSI initiator-specific target.
Port WWN(s) Secondary	The Fibre Channel target's port addresses associated with this iSCSI initiator-specific target.
LUN Map (Hex) iSCSI	The Fibre Channel target's port addresses associated with this iSCSI initiator-specific target.
LUN Map (Hex) FC Primary	The Fibre Channel target's port addresses associated with this iSCSI initiator-specific target.
LUN Map (Hex) FC Secondary	The Fibre Channel target's port addresses associated with this iSCSI initiator-specific target.

Field	Description
No AutoZone Creation	Indicates if a Fibre Channel zone is automatically created for this iSCSI initiator-target and the iSCSI initiator. If true the zone is not automatically created. If false (default) the zone is automatically created.
Trespass Mode	The trespass mode for this node. If true the Fibre Channel node instance presents all LUN I/O requests to the secondary port (fcSecondaryAddress) if the primary port (fcAddress) is down.
Revert to Primary Port	The revert to primary mode for this node. If true the Fibre Channel node instance presents all LUN I/O requests to the primary port (fcAddress) when the primary port comes back online.
Primary PWWN VSAN	Indicates the VSAN into which the auto zone is placed for this initiator target. If this object is not set then the VSAN is determined by querying the name server.
Secondary PWWN VSAN	Indicates the VSAN into which the auto zone is placed for this initiator target. If this object is not set then the VSAN is determined by querying the name server.

iSCSI Initiator PWWN

Field	Description
Port WWN	The Fibre Channel address for this entry.

iSCSI Sessions

Field	Description
Type	Type of iSCSI session: <ul style="list-style-type: none"> • normal—session is a normal iSCSI session • discovery—session is being used only for discovery.
TargetName	If Direction is Outbound, this will contain the name of the remote target.
Vsan ID	The VSAN to which this session belongs to.
ISID	The initiator-defined portion of the iSCSI session ID.
TSIH	The target-defined identification handle for this session.

iSCSI Sessions Detail

Field	Description
ConnectionNumber	The number of transport protocol connections that currently belong to this session.

Field	Description
ImmediateData	Whether the initiator and target have agreed to support immediate data on this session.
Initial	If true, the initiator must wait for a Ready-To-Transfer before sending to the target. If false, the initiator may send data immediately, within limits set by FirstBurstSize and the expected data transfer length of the request.
MaxOutstanding	The maximum number of outstanding Ready-To-Transfers per task within this session.
First	The maximum length supported for unsolicited data sent within this session.
Max	The maximum number of bytes which can be sent within a single sequence of Data-In or Data-Out PDUs.
Sequence	If false, indicates that iSCSI data PDU sequences may be transferred in any order. If true indicates that data PDU sequences must be transferred using continuously increasing offsets, except during error recovery.
PDU	If false, iSCSI data PDUs within sequences may be in any order. If true indicates that data PDUs within sequences must be at continuously increasing addresses, with no gaps or overlay between PDUs.

Additional References

For additional information related to implementing IP storage, see the following section:

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	Cisco MDS 9000 Family Command Reference, Release 5.0(1a)

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified.	To locate and download MIBs, go to the following URL: http://www.cisco.com/dc-os/mibs



CHAPTER 43

Configuring IPv4 for Gigabit Ethernet Interfaces

- [Configuring IPv4 for Gigabit Ethernet Interfaces](#) , on page 927

Configuring IPv4 for Gigabit Ethernet Interfaces

Cisco MDS 9000 Family switches support IP version 4 (IPv4) on Gigabit Ethernet interfaces. This chapter describes how to configure IPv4 addresses and other IPv4 features.

This chapter includes the following topics:

Information About IPv4

Cisco MDS 9000 Family supports IP version 4 (IPv4) on Gigabit Ethernet interfaces. Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each IPS module or MPS-14/2 module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured.

A new port mode, called IPS, is defined for Gigabit Ethernet ports on each IPS module or MPS-14/2 module. IP storage ports are implicitly set to IPS mode, so it can only be used to perform iSCSI and FCIP storage functions. IP storage ports do not bridge Ethernet frames or route other IP packets.

Each IPS port represents a single virtual Fibre Channel host in the Fibre Channel SAN. All the iSCSI hosts connected to this IPS port are merged and multiplexed through the single Fibre Channel host.



Note The Gigabit Ethernet interfaces on the MPS-14/2 module do not support EtherChannel.

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each IPS module or MPS-14/2 module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured. This section covers the steps required to configure IP for subsequent use by FCIP and iSCSI.



Note For information about configuring FCIP, see *Chapter 38, “Configuring FCIP”*. For information about configuring iSCSI, see *Chapter 40, “Configuring iSCSI”*.

A new port mode, called IPS, is defined for Gigabit Ethernet ports on each IPS module or MPS-14/2 module. IP storage ports are implicitly set to IPS mode, so it can only be used to perform iSCSI and FCIP storage functions. IP storage ports do not bridge Ethernet frames or route other IP packets.

Each IPS port represents a single virtual Fibre Channel host in the Fibre Channel SAN. All the iSCSI hosts connected to this IPS port are merged and multiplexed through the single Fibre Channel host.

In large scale iSCSI deployments where the Fibre Channel storage subsystems do not require explicit LUN access control for every host device, use of proxy-initiator mode simplifies the configuration.



Note The Gigabit Ethernet interfaces on the MPS-14/2 module do not support EtherChannel.



Note To configure IPv6 on a Gigabit Ethernet interface, see the “*Configuring IPv6 Addressing and Enabling IPv6 Routing*” section .



Tip Gigabit Ethernet ports on any IPS module or MPS-14/2 module should not be configured in the same Ethernet broadcast domain as the management Ethernet port. They should be configured in a different broadcast domain, either by using separate standalone hubs or switches or by using separate VLANs.

This section includes the following topics:

Interface Descriptions

See the Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide for details on configuring the switch port description for any interface.

Beacon Mode

See the Interfaces Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide for details on configuring the beacon mode for any interface.

About VLANs for Gigabit Ethernet

Virtual LANs (VLANs) create multiple virtual Layer 2 networks over a physical LAN network. VLANs provide traffic isolation, security, and broadcast control.

Gigabit Ethernet ports automatically recognize Ethernet frames with IEEE 802.1Q VLAN encapsulation. If you need to have traffic from multiple VLANs terminated on one Gigabit Ethernet port, configure subinterfaces—one for each VLAN.

If the IPS module or MPS-14/2 module is connected to a Cisco Ethernet switch, and you need to have traffic from multiple VLANs coming to one IPS port, verify the following requirements on the Ethernet switch:

- The Ethernet switch port connected to the IPS module or MPS-14/2 module is configured as a trunking port.
- The encapsulation is set to 802.1Q and not ISL, which is the default.

Use the VLAN ID as a subscription to the Gigabit Ethernet interface name to create the subinterface name:

slot-number / port-number.VLAN-ID

Interface Subnet Requirements

Gigabit Ethernet interfaces (major), subinterfaces (VLAN ID), and management interfaces (mgmt 0) can be configured in the same or different subnet depending on the configuration (see [Table 112: Subnet Requirements for Interfaces](#), on page 929).

Table 112: Subnet Requirements for Interfaces

Interface 1	Interface 2	Same Subnet Allowed	Notes
Gigabit Ethernet 1/1	Gigabit Ethernet 1/2	Yes	Two major interfaces can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.100	Yes	Two subinterfaces with the same VLAN ID can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.200	No	Two subinterfaces with different VLAN IDs cannot be configured in the same subnet.
Gigabit Ethernet 1/1	Gigabit Ethernet 1/1.100	No	A subinterface cannot be configured on the same subnet as the major interface.
mgmt0	Gigabit Ethernet 1/1.100	No	The mgmt0 interface cannot be configured in the same subnet as the Gigabit Ethernet interfaces or subinterfaces.
mgmt0	Gigabit Ethernet 1/1	No	



Note The configuration requirements in [Table 112: Subnet Requirements for Interfaces](#), on page 929 also apply to Ethernet PortChannels.

Licensing Requirements for IPv4 for Gigabit Ethernet Interfaces

The following table shows the licensing requirements for this feature:

License	License Description
Enterprise package (ENTERPRISE_PKG)	It comprises IPsec and IKE for IPv4.

Guidelines and Limitations

Follow these guidelines when configuring IPv4-ACLs for Gigabit Ethernet interfaces:

- Only use Transmission Control Protocol (TCP) or Internet Control Message Protocol (ICMP).

**Note**

Other protocols such as User Datagram Protocol (UDP) and HTTP are not supported in Gigabit Ethernet interfaces. Applying an ACL that contains rules for these protocols to a Gigabit Ethernet interface is allowed but those rules have no effect.

- Apply IPv4-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
 - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
 - The **established** option is ignored when you apply IPv4-ACLs containing this option to Gigabit Ethernet interfaces.
 - If an IPv4-ACL rule applies to a pre-existing TCP connection, that rule is ignored. For example if there is an existing TCP connection between A and B and an IPv4-ACL which specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.

**Tip**

If IPv4-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group.

Default Settings

[Table 113: Default IPv4 Parameters](#) , on page 930 lists the default settings for IPv4 parameters.

Table 113: Default IPv4 Parameters

Parameters	Default
IPv4 MTU frame size	1500 bytes for all Ethernet ports
Autonegotiation	Enabled
Promiscuous mode	Disabled

Configuring IPv4

This section includes the following topics:

Configuring Gigabit Ethernet Interface

To configure the Gigabit Ethernet interface, follow these steps:

Procedure

Step 1

Expand Switches > Interfaces > Ethernet > IPS.

You see the Gigabit Ethernet Configuration in the Information pane.

- Step 2** Click the IP Addresses tab.
- Step 3** Click Create Row.
You see the Create Gigabit Ethernet Interface dialog box.
- Step 4** Select the switch on which you want to create the Gigabit Ethernet interface.
- Step 5** Enter the interface. For example, 2/2 for slot 2, port 2.
- Step 6** Enter the IPv4 address (10.1.1.100) and subnet mask (255.255.255.0).
- Step 7** Click Create to save these changes or click Close to discard any unsaved changes.
-

Configuring Autonegotiation

By default, autonegotiation is enabled all Gigabit Ethernet interface. You can enable or disable autonegotiation for a specified Gigabit Ethernet interface. When autonegotiation is enabled, the port automatically detects the speed or pause method, and duplex of incoming signals based on the link partner. You can also detect link up conditions using the autonegotiation feature.

To configure autonegotiation, follow these steps:

Procedure

- Step 1** Expand Switches > Interfaces > Ethernet > IPS.
You see the Gigabit Ethernet Configuration in the Information pane.
- Step 2** In the General tab, you can enable or disable the Auto Negotiate option for a specific switch.
- Step 3** Click **Apply Changes**.
-

Configuring the MTU Frame Size

You can configure the interfaces on a switch to transfer large (or jumbo) frames on a port. The default IP maximum transmission unit (MTU) frame size is 1500 bytes for all Ethernet ports. By configuring jumbo frames on a port, the MTU size can be increased up to 9000 bytes.



Note The minimum MTU size is 576 bytes.



Tip MTU changes are disruptive, all FCIP links and iSCSI sessions flap when the software detects a change in the MTU size.

You do not need to explicitly issue the **shutdown** and **no shutdown** commands.

To configure the MTU frame size, follow these steps:

Procedure

- Step 1** Expand Switches > Interfaces > Ethernet > IPS.
You see the Gigabit Ethernet Configuration in the Information pane.
- Step 2** In the General tab, in the Mtu column, you can enter a new value to configure the MTU Frame Size for a specific switch. For example 3000 bytes. The default is 1500 bytes.
- Step 3** Click **Apply Changes**.
-

Configuring Promiscuous Mode

You can enable or disable promiscuous mode on a specific Gigabit Ethernet interface. By enabling the promiscuous mode, the Gigabit Ethernet interface receives all the packets and the software then filters and discards the packets that are not destined for that Gigabit Ethernet interface.

To configure the promiscuous mode, follow these steps:

Procedure

- Step 1** Expand Switches > Interfaces > Ethernet > IPS.
You see the Gigabit Ethernet Configuration in the Information pane.
- Step 2** In the General tab, you can enable or disable the Promiscuous Mode option for a specific switch.
- Step 3** Click **Apply Changes**.
-

Configuring the VLAN Subinterface

To configure a VLAN subinterface (VLAN ID), follow these steps:

Procedure

- Step 1** Select Interface > Ethernet and iSCSI.
- Step 2** Click the Sub Interfaces tab.
- Step 3** Select the Gigabit Ethernet subinterface on which 802.1Q should be used.
- Step 4** Click the Edit IP Address button.
- Step 5** Enter the IPv4 address and subnet mask for the Gigabit Ethernet interface.
- Step 6** Click **Create** to save the changes or you may click **Close**.
-

Additional References

For additional information related to implementing FCIPs, see the following section:

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	Cisco MDS 9000 Family Command Reference, Release 5.0(1a)

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
CISCO-IETF-VRRP-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/dc-os/mibs



CHAPTER 44

Configuring IPv6 for Gigabit Ethernet Interfaces

- [Configuring IPv6 for Gigabit Ethernet Interfaces, on page 935](#)

Configuring IPv6 for Gigabit Ethernet Interfaces

IP version 6 (IPv6) provides extended addressing capability beyond those provided in IP version 4 (IPv4) in Cisco MDS NX-OS. The architecture of IPv6 has been designed to allow existing IPv4 users to transition easily to IPv6 while providing services such as end-to-end security, quality of service (QoS), and globally unique addresses.



Note

For Cisco NX-OS features that use IP addressing, refer to the chapters in this guide that describe those features for information on IPv6 addressing support.



Note

To configure IP version 4 (IPv4) on a Gigabit Ethernet interface, see [Chapter 43, “Configuring IPv4 for Gigabit Ethernet Interfaces.”](#)

This chapter includes the following topics:

Information About IPV6

IP version 6 (IPv6) provides extended addressing capability beyond those provided in IP version 4 (IPv4) in Cisco MDS NX-OS by quadrupling the number of network address bits from 32 bits (in IPv4) to 128 bits. The architecture of IPv6 has been designed to allow existing IPv4 users to transition easily to IPv6 while providing services such as end-to-end security, quality of service (QoS), and globally unique addresses.

IPv6 provides the following enhancements over IPv4:

- Allows networks to scale and provide global reachability.
- Reduces the need for private address and network address translation (NAT).
- Provides simpler autoconfiguration of addresses.

This section describes the IPv6 features supported by Cisco MDS NX-OS and includes the following topics:

Extended IPv6 Address Space for Unique Addresses

IPv6 extends the address space by quadrupling the number of network address bits from 32 bits (in IPv4) to 128 bits, which provides many more globally unique IP addresses. By being globally unique, IPv6 addresses enable global reachability and end-to-end security for networked devices, functionality that is crucial to the applications and services that are driving the demand for more addresses.

IPv6 Address Formats

IPv6 addresses are represented as a series of 16-bit hexadecimal fields separated by colons (:) in the format x:x:x:x:x:x:x:x. The following are examples of IPv6 addresses:

2001:0DB8:7654:3210:FEDC:BA98:7654:3210

2001:0DB8:0:0:8:800:200C:417A

It is common for IPv6 addresses to contain successive hexadecimal fields of zeros. To make IPv6 addresses easier to use, two colons (::) may be used to compress successive hexadecimal fields of zeros at the beginning, middle, or end of an IPv6 address (the colons represent successive hexadecimal fields of zeros). [Table 114: Compressed IPv6 Address Formats](#) , on page 936 lists compressed IPv6 address formats.



Note Two colons (::) can be used only once in an IPv6 address to represent the longest successive hexadecimal fields of zeros.



Note The hexadecimal letters in IPv6 addresses are not case-sensitive.

Table 114: Compressed IPv6 Address Formats

IPv6 Address Type	Uncompressed Format	Compressed Format
Unicast	2001:0DB8:800:200C:0:0:0:417A	2001:0DB8:800:200C::417A
Multicast	FF01:0:0:0:0:0:0:101	FF01::101

IPv6 Address Prefix Format

An IPv6 address prefix, in the format *ipv6-prefix/prefix-length* , can be used to represent bit-wise contiguous blocks of the entire address space. The *ipv6-prefix* is specified in hexadecimal using 16-bit values between the colons. The *prefix-length* is a decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). For example, 2001:0DB8:8086:6502::/32 is a valid IPv6 prefix.

IPv6 Address Type: Unicast


An IPv6 unicast address is an identifier for a single interface on a single node. A packet that is sent to a unicast address is delivered to the interface identified by that address. The Cisco MDS NX-OS supports the following IPv6 unicast address types:

- Global addresses
- Link-local addresses

Global Addresses

Global IPv6 addresses are defined by a global routing prefix, a subnet ID, and an interface ID. [Figure 154: Global Address Format, on page 937](#) shows the structure of a global address.

Figure 154: Global Address Format



Addresses with a prefix of 2000::/3 (001) through E000::/3 (111) are required to have 64-bit interface identifiers in the extended universal identifier (EUI)-64 format. The Internet Assigned Numbers Authority (IANA) allocates the IPv6 address space in the range of 2000::/16 to regional registries.


The aggregatable global address typically consists of a 48-bit global routing prefix and a 16-bit subnet ID or Site-Level Aggregator (SLA). In the IPv6 aggregatable global unicast address format document (RFC 2374), the global routing prefix included two other hierarchically structured fields named Top-Level Aggregator (TLA) and Next-Level Aggregator (NLA). The IETF decided to remove the TLA and NLA fields from the RFCs because these fields are policy-based. Some existing IPv6 networks deployed before the change might still be using networks based on the older architecture.

A 16-bit subnet field called the subnet ID could be used by individual organizations to create their own local addressing hierarchy and to identify subnets. A subnet ID is similar to a subnet in IPv4, except that an organization with an IPv6 subnet ID can support up to 65,535 individual subnets.

An interface ID is used to identify interfaces on a link. The interface ID must be unique to the link. They may also be unique over a broader scope. In many cases, an interface ID will be the same as, or based on, the link-layer address of an interface, which results in a globally unique interface ID. Interface IDs used in aggregatable global unicast and other IPv6 address types must be 64 bits long and constructed in the modified EUI-64 format.

Cisco MDS NX-OS supports IEEE 802 interface types (for example, Gigabit Ethernet interfaces). The first three octets (24 bits) are taken from the Organizationally Unique Identifier (OUI) of the 48-bit link-layer address (MAC address) of the interface, the fourth and fifth octets (16 bits) are a fixed hexadecimal value of FFFE, and the last three octets (24 bits) are taken from the last three octets of the MAC address. The construction of the interface ID is completed by setting the Universal/Local (U/L) bit—the seventh bit of the first octet—to a value of 0 or 1. A value of 0 indicates a locally administered identifier; a value of 1 indicates a globally unique IPv6 interface identifier (see [Figure 155: Interface Identifier Format, on page 938](#)).

Figure 155: Interface Identifier Format




Link-Local Address

A link-local address is an IPv6 unicast address that is automatically configured on an interface using the link-local prefix FE80::/10 and the interface identifier in the modified EUI-64 format. Link-local addresses are used in the neighbor discovery protocol and the stateless autoconfiguration process. Nodes on a local link can use link-local addresses to communicate.

Figure 156: Link-Local Address Format, on page 938 shows the structure of a link-local address.

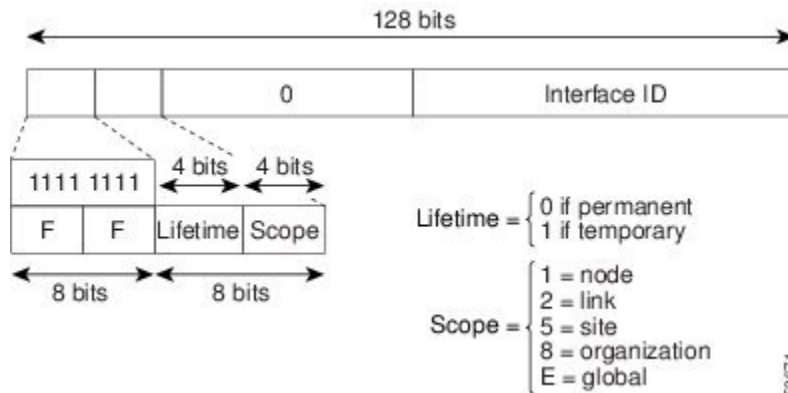
Figure 156: Link-Local Address Format



IPv6 Address Type: Multicast

An IPv6 multicast address is an IPv6 address that has a prefix of FF00::/8 (1111 1111). An IPv6 multicast address is an identifier for a set of interfaces that typically belong to different nodes. A packet sent to a multicast address is delivered to all interfaces identified by the multicast address. The second octet following the prefix defines the lifetime and scope of the multicast address. A permanent multicast address has a lifetime parameter equal to 0; a temporary multicast address has a lifetime parameter equal to 1. A multicast address has the scope of a node, link, site, or organization, or a global scope has a scope parameter of 1, 2, 5, 8, or E, respectively. For example, a multicast address with the prefix FF02::/16 is a permanent multicast address with a link scope. Figure 157: IPv6 Multicast Address Format, on page 939 shows the format of the IPv6 multicast address.

Figure 157: IPv6 Multicast Address Format




IPv6 hosts are required to join (receive packets destined for) the following multicast groups:

- All-node multicast group FF02::1.
- Solicited-node multicast group FF02:0:0:0:0:1:FF00:0000/104 concatenated with the low-order 24 bit of the unicast address.

The solicited-node multicast address is a multicast group that corresponds to an IPv6 unicast address. IPv6 nodes must join the associated solicited-node multicast group for every unicast address to which it is assigned. The IPv6 solicited-node multicast address has the prefix FF02:0:0:0:0:1:FF00:0000/104 concatenated with the 24 low-order bits of a corresponding IPv6 unicast address. (See [Figure 158: IPv6 Solicited-Node Multicast Address Format, on page 939](#)) For example, the solicited-node multicast address corresponding to the IPv6 address 2037::01:800:200E:8C6C is FF02::1:FF0E:8C6C. Solicited-node addresses are used in neighbor solicitation messages.

Figure 158: IPv6 Solicited-Node Multicast Address Format

**Note**

There are no broadcast addresses in IPv6. IPv6 multicast addresses are used instead of broadcast addresses.

ICMP for IPv6


Internet Control Message Protocol (ICMP) in IPv6 functions the same as ICMP in IPv4—ICMP generates error messages such as ICMP destination unreachable messages, and informational messages such as ICMP echo request and reply messages. Additionally, ICMP packets in IPv6 are used in the IPv6 neighbor discovery

process, path MTU discovery, and the Multicast Listener Discovery (MLD) protocol for IPv6. MLD is based on version 2 of the Internet Group Management Protocol (IGMP) for IPv4.

A value of 58 in the Next Header field of the basic IPv6 packet header identifies an IPv6 ICMP packet. ICMP packets in IPv6 resemble a transport-layer packet in the sense that the ICMP packet follows all the extension headers and is the last piece of information in the IPv6 packet. Within IPv6 ICMP packets, the ICMPv6 Type and ICMPv6 Code fields identify IPv6 ICMP packet specifics, such as the ICMP message type. The value in the Checksum field is derived (computed by the sender and checked by the receiver) from the fields in the IPv6 ICMP packet and the IPv6 pseudoheader. The ICMPv6 Data field contains error or diagnostic information relevant to IP packet processing.

Figure 159: IPv6 ICMP Packet Header Format, on page 940 shows the IPv6 ICMP packet header format.

Figure 159: IPv6 ICMP Packet Header Format



Path MTU Discovery for IPv6

As in IPv4, path MTU discovery in IPv6 allows a host to dynamically discover and adjust to differences in the MTU size of every link along a given data path. In IPv6, however, fragmentation is handled by the source of a packet when the path MTU of one link along a given data path is not large enough to accommodate the size of the packets. Having IPv6 hosts handle packet fragmentation saves IPv6 router processing resources and helps IPv6 networks run more efficiently.



Note

In IPv4, the minimum link MTU is 68 octets, which means that the MTU size of every link along a given data path must support an MTU size of at least 68 octets. In IPv6, the minimum link MTU is 1280 octets. We recommend using MTU value of 1500 octets for IPv6 links.

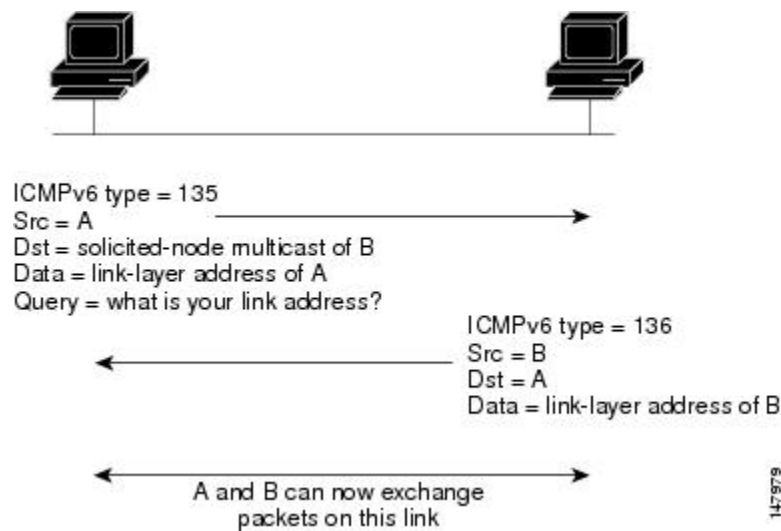
IPv6 Neighbor Discovery

The IPv6 neighbor discovery process uses ICMP messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), verify the reachability of a neighbor, and keep track of neighboring routers.

IPv6 Neighbor Solicitation and Advertisement Messages

A value of 135 in the Type field of the ICMP packet header identifies a neighbor solicitation message. Neighbor solicitation messages are sent on the local link when a node wants to determine the link-layer address of another node on the same local link. (See [Figure 160: IPv6 Neighbor Discovery—Neighbor Solicitation Message, on page 941](#).) When a node wants to determine the link-layer address of another node, the source address in a neighbor solicitation message is the IPv6 address of the node sending the neighbor solicitation message. The destination address in the neighbor solicitation message is the solicited-node multicast address that corresponds to the IPv6 address of the destination node. The neighbor solicitation message also includes the link-layer address of the source node.

Figure 160: IPv6 Neighbor Discovery—Neighbor Solicitation Message



After receiving the neighbor solicitation message, the destination node replies by sending a neighbor advertisement message, which has a value of 136 in the Type field of the ICMP packet header, on the local link. The source address in the neighbor advertisement message is the IPv6 address of the node (more specifically, the IPv6 address of the node interface) sending the neighbor advertisement message. The destination address in the neighbor advertisement message is the IPv6 address of the node that sent the neighbor solicitation message. The data portion of the neighbor advertisement message includes the link-layer address of the node sending the neighbor advertisement message.

After the source node receives the neighbor advertisement, the source node and destination node can communicate.

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. When a node wants to verify the reachability of a neighbor, the destination address in a neighbor solicitation message is the unicast address of the neighbor.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link. When there is such a change, the destination address for the neighbor advertisement is the all-node multicast address.

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. Neighbor unreachability detection identifies the failure of a neighbor or the failure of the forward path to the neighbor, and is used for all paths between hosts and neighboring nodes (hosts or routers). Neighbor unreachability detection is performed for neighbors to which only unicast packets are being sent and is not performed for neighbors to which multicast packets are being sent.

A neighbor is considered reachable when the neighbor returns a positive acknowledgment indicating that it has received and processed packets previously sent to it. A positive acknowledgment could be from an upper-layer protocol such as TCP indicating that a connection is making forward progress (reaching its destination) or the receipt of a neighbor advertisement message in response to a neighbor solicitation message. If packets are reaching the peer, they are also reaching the next-hop neighbor of the source. Therefore, forward progress is also a confirmation that the next-hop neighbor is reachable.

For destinations that are not on the local link, forward progress implies that the first-hop router is reachable. When acknowledgments from an upper-layer protocol are not available, a node probes the neighbor using unicast neighbor solicitation messages to verify that the forward path is still working. The return of a solicited neighbor advertisement message from the neighbor is a positive acknowledgment that the forward path is still working (neighbor advertisement messages that have the solicited flag set to a value of 1 are sent only in response to a neighbor solicitation message). Unsolicited messages confirm only the one-way path from the source to the destination node; solicited neighbor advertisement messages indicate that a path is working in both directions.

**Note**

A neighbor advertisement message that has the solicited flag set to a value of 0 must not be considered as a positive acknowledgment that the forward path is still working.

Neighbor solicitation messages are also used in the stateless autoconfiguration process to verify the uniqueness of unicast IPv6 addresses before the addresses are assigned to an interface. Duplicate address detection is performed first on a new, link-local IPv6 address before the address is assigned to an interface (the new address remains in a tentative state while duplicate address detection is performed). Specifically, a node sends a neighbor solicitation message with an unspecified source address and a tentative link-local address in the body of the message. If another node is already using that address, the node returns a neighbor advertisement message that contains the tentative link-local address. If another node is simultaneously verifying the uniqueness of the same address, that node also returns a neighbor solicitation message. If no neighbor advertisement messages are received in response to the neighbor solicitation message and no neighbor solicitation messages are received from other nodes that are attempting to verify the same tentative address, the node that sent the original neighbor solicitation message considers the tentative link-local address to be unique and assigns the address to the interface.

Every IPv6 unicast address (global or link-local) must be checked for uniqueness on the link; however, until the uniqueness of the link-local address is verified, duplicate address detection is not performed on any other IPv6 addresses associated with the link-local address.

Router Discovery

Router discovery performs both router solicitation and router advertisement. Router solicitations are sent by hosts to all-routers multicast addresses. Router advertisements are sent by routers in response to solicitations or unsolicited and contain default router information as well as additional parameters such as the MTU and hop limit.


IPv6 Stateless Autoconfiguration

All interfaces on IPv6 nodes must have a link-local address, which is automatically configured from the identifier for an interface and the link-local prefix FE80::/10. A link-local address enables a node to communicate with other nodes on the link and can be used to further configure the node.

Nodes can connect to a network and automatically generate site-local and global IPv6 address without the need for manual configuration or help of a server, such as a DHCP server. With IPv6, a router on the link

advertises in router advertisement (RA) messages any site-local and global prefixes, and its willingness to function as a default router for the link. RA messages are sent periodically and in response to router solicitation messages, which are sent by hosts at system startup. (See [Figure 161: IPv6 Stateless Autoconfiguration](#), on page 943.)

Figure 161: IPv6 Stateless Autoconfiguration

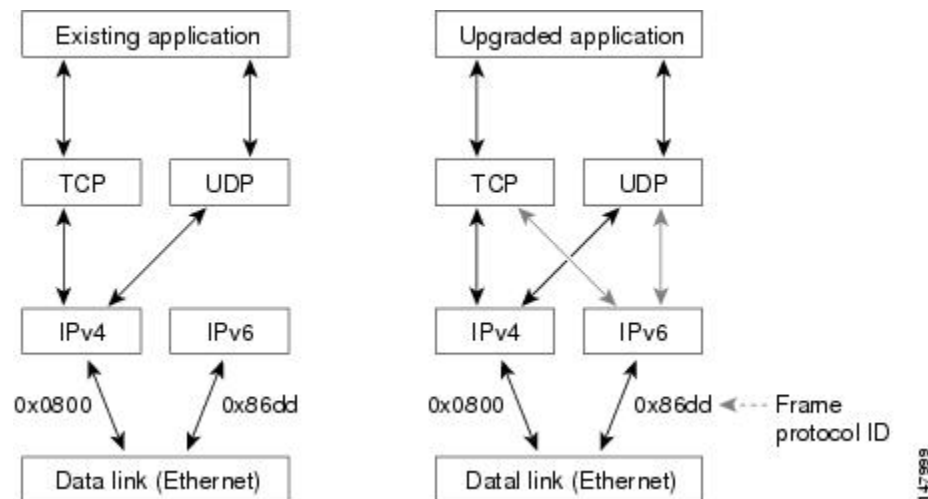


A node on the link can automatically configure site-local and global IPv6 addresses by appending its interface identifier (64 bits) to the prefixes (64 bits) included in the RA messages. The resulting 128-bit IPv6 addresses configured by the node are then subjected to duplicate address detection to ensure their uniqueness on the link. If the prefixes advertised in the RA messages are globally unique, then the IPv6 addresses configured by the node are also guaranteed to be globally unique. Router solicitation messages, which have a value of 133 in the Type field of the ICMP packet header, are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled RA message.

Dual IPv4 and IPv6 Protocol Stacks

The dual IPv4 and IPv6 protocol stack technique is one technique for a transition to IPv6. It enables gradual, one-by-one upgrades to applications running on nodes. Applications running on nodes are upgraded to make use of the IPv6 protocol stack. Applications that are not upgraded—they support only the IPv4 protocol stack—can coexist with upgraded applications on the same node. New and upgraded applications simply make use of both the IPv4 and IPv6 protocol stacks. (See [Figure 162: Dual IPv4 and IPv6 Protocol Stack Technique](#), on page 943.)

Figure 162: Dual IPv4 and IPv6 Protocol Stack Technique




A new API has been defined to support both IPv4 and IPv6 addresses and DNS requests. An application can be upgraded to the new API and still use only the IPv4 protocol stack. The Cisco MDS NX-OS supports the

dual IPv4 and IPv6 protocol stack technique. When an interface is configured with both an IPv4 and an IPv6 address, the interface will accept and process both IPv4 and IPv6 traffic.

In [Figure 163: Dual IPv4 and IPv6 Protocol Stack Applications, on page 944](#), an application that supports dual IPv4 and IPv6 protocol stacks requests all available addresses for the destination host name `www.a.com` from a DNS server. The DNS server replies with all available addresses (both IPv4 and IPv6 addresses) for `www.a.com`. The application chooses an address—in most cases, IPv6 addresses are the default choice—and connects the source node to the destination using the IPv6 protocol stack.

Figure 163: Dual IPv4 and IPv6 Protocol Stack Applications



IPv6 Addressing and Enabling IPv6 Routing

IPv6 addresses are represented as a series of 16-bit hexadecimal fields separated by colons (:) in the format `x:x:x:x:x:x:x:x`. It is common for IPv6 addresses to contain successive hexadecimal fields of zeros. To make IPv6 addresses easier to use, two colons (:) may be used to compress successive hexadecimal fields of zeros at the beginning, middle, or end of an IPv6 address (the colons represent successive hexadecimal fields of zeros). By default, IPv6 addresses are not configured, and IPv6 processing is disabled. You can configure IPv6 addresses on the following interface types:

- Gigabit Ethernet
- Management
- VLAN (Gigabit Ethernet subinterface)
- VSAN



Note

The IPv6 address *ipv6-address* argument in the **ipv6 address** command must be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons.

The IPv6 prefix *ipv6-prefix* argument in the **ipv6 address** command must be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons.

The IPv6 prefix length *prefix-length* argument in the **ipv6 address** command is a decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). A slash mark must precede the decimal value.

Configuring a global IPv6 address on an interface automatically configures a link-local address and activates IPv6 for that interface. Additionally, the configured interface automatically joins the following required multicast groups for that link:

- Solicited-node multicast group `FF02:0:0:0:0:1:FF00::/104` for each unicast address assigned to the interface

- All-node link-local multicast group FF02::1

This task explains how to assign IPv6 addresses to individual router interfaces and enable the processing of IPv6 traffic. By default, IPv6 addresses are not configured and IPv6 processing is disabled.

You can configure IPv6 addresses on the following interface types:

- Gigabit Ethernet
- Management
- VLAN (Gigabit Ethernet subinterface)
- VSAN

Configuring a global IPv6 address on an interface automatically configures a link-local address and activates IPv6 for that interface. Additionally, the configured interface automatically joins the following required multicast groups for that link:

- Solicited-node multicast group FF02:0:0:0:1:FF00::/104 for each unicast address assigned to the interface
- All-node link-local multicast group FF02::1



Note The solicited-node multicast address is used in the neighbor discovery process.



Note The maximum number of IPv6 addresses (static and autoconfigured) allowed on an interface is eight, except on the management (mgmt 0) interface where only one static IPv6 address can be configured.

Transitioning from IPv4 to IPv6

Cisco MDS NX-OS does not support any transitioning mechanisms from IPv4 to IPv6. However, you can use the transitioning schemes in the Cisco router products for this purpose. For information on configuring Cisco routers to transition your network, refer to the “Implementing Tunneling for IPv6” chapter in the [Cisco IOS IPv6 Configuration Guide](#).

Guidelines and Limitations



Tip If IPv6-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to a Ethernet PortChannel group. See the Security Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Security Configuration Guide for information on configuring IPv6-ACLs.

Follow these guidelines when configuring IPv6-ACLs for Gigabit Ethernet interfaces:

- Only use Transmission Control Protocol (TCP) or Internet Control Message Protocol (ICMP).



Note Other protocols such as User Datagram Protocol (UDP) and HTTP are not supported in Gigabit Ethernet interfaces. Applying an ACL that contains rules for these protocols to a Gigabit Ethernet interface is allowed but those rules have no effect.

- Apply IPv6-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
 - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
 - The **established** option is ignored when you apply IPv6-ACLs containing this option to Gigabit Ethernet interfaces.
 - If an IPv6-ACL rule applies to a preexisting TCP connection, that rule is ignored. For example, if there is an existing TCP connection between A and B and an IPv6-ACL that specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.

See the Security Configuration Guide, Cisco DCNM for SANCisco MDS 9000 Family NX-OS Security Configuration Guide for information on applying IPv6-ACLs to an interface.

Default Settings

[Table 115: Default IPv6 Parameters](#), on page 946 lists the default settings for IPv6 parameters.

Table 115: Default IPv6 Parameters

Parameters	Default
IPv6 processing	Disabled
Duplicate address detection attempts	0 (neighbor discovery disabled)
Reachability time	1000 milliseconds
Retransmission time	30000 milliseconds
IPv6-ACLs	None

Configuring Basic Connectivity for IPv6

This section includes the following topics:

Configuring IPv6 Addressing and Enabling IPv6 Routing

To configure an IPv6 address on an interface using Device Manager, follow these steps:

Procedure

-
- Step 1** Choose Interfaces > Gigabit Ethernet and iSCSI.
You see the Gigabit Ethernet Configuration dialog box.
- Step 2** Click the IP Address that you want to configure and click Edit IP Address.
You see the IP Address dialog box.

- Step 3** Click Create and set the IP Address/Mask field, using the IPv6 format (for example, 2001:0DB8:800:200C::417A/64).
- Step 4** Click Create to save these changes or click Close to discard any unsaved changes.
-

Configuring IPv6 Routing using Device Manager

To enable IPv6 routing using Device Manager, follow these steps:

Procedure

- Step 1** Choose IP > Routing. You see the IP Routing Configuration dialog box.
- Step 2** Check the Routing Enabled check box.
- Step 3** Click Apply to save these changes or click Close to discard any unsaved changes.
-

Configuring IPv4 and IPv6 Protocol Addresses

When an interface in a Cisco networking device is configured with both an IPv4 and an IPv6 address, the interface can send and receive data on both IPv4 and IPv6 networks.

To configure an interface in a Cisco networking device to support both the IPv4 and IPv6 protocol stacks using Device Manager, follow these steps:

Procedure

- Step 1** Choose Interfaces > Gigabit Ethernet and iSCSI.
You see the Gigabit Ethernet Configuration dialog box.
- Step 2** Click the IP Address field that you want to configure and click Edit IP Address.
You see the IP Address dialog box.
- Step 3** Click Create and set the IP Address/Mask field, using the IPv4 or IPv6 format.
- Step 4** Click Create to save these changes or click Close to discard any unsaved changes.
-

Configuring Neighbor Discovery Parameters

You can configure the following neighbor discovery parameters:

- Duplicate address detection attempts
- Reachability time
- Retransmission timer



Note We recommend that you use the factory-defined defaults for these parameters.

This section includes the following topics:

Configuring a IPv6 Static Route

You must manually configure IPv6 static routes and define an explicit path between two networking devices. IPv6 static routes are not automatically updated and must be manually reconfigured if the network topology changes.

To configure a IPv6 static route using Device Manager, follow these steps:

Procedure

- Step 1** Choose IP > Routing.
You see the IP Routing Configuration dialog box.
 - Step 2** Click Create.
You see the Create IP Route dialog box.
 - Step 3** Set the Dest field to the IPv6 destination address.
 - Step 4** Set the Mask field to the IPv6 subnet mask.
 - Step 5** Set the Gateway field to the IPv6 default gateway.
 - Step 6** (Optional) Set the Metric field to the desired route metric.
 - Step 7** Select the interface from the Interface drop-down menu.
 - Step 8** Click Create to save these changes or click Close to discard any unsaved changes.
-



CHAPTER 45

Configuring SCSI Flow Services

- [Configuring SCSI Flow Services, on page 949](#)

Configuring SCSI Flow Services

This chapter describes SCSI flow services which is supported on the Storage Services Module (SSM).

This chapter includes the following sections:

Information About SCSI Flow Services

This section includes the following topics:


SCSI Flow Services Overview

An SCSI initiator and target combination is an SCSI flow. SCSI flow services provide enhanced features for SCSI flows, such as Write Acceleration and flow monitoring for statistics obtained on an SSM.

The SCSI flow services functional architecture consists of the following components:

- SCSI flow manager (SFM) on the supervisor: The SFM resides on a supervisor module and handles the configuration of SCSI flows, validating them and relaying configuration information to the appropriate SSM. It also handles any dynamic changes to the status of the SCSI flow due to external events and registers changes that occur due to various operations.
- SCSI flow configuration CLI on the supervisor: The SFCC resides on the CPP of the SSM. It receives flow configuration requests from the SFM, programs the DPP corresponding to the initiator and target port interfaces, and responds to the SFM with the status of the configuration request.
- SCSI flow configuration client on the Control Path Processor (CPP) of an SSM.
- SCSI flow feature set support on the Data Path Processor (DPP) of an SSM: The DPP on the SSM examines all the messages between the initiator and target and provides SCSI flow features, such as Fibre Channel Write Acceleration and statistics monitoring.

The following figure shows an example of the SCSI flow services functional architecture.



Note The SCSI target and initiator must be connected to different SSMs on different switches.



Note For statistics monitoring, the target device is not required to be connected to an SSM.

SCSI Flow Specification Attributes

A SCSI flow specification consists of the following attributes:

- SCSI flow identifier
- VSAN identifier
- SCSI initiator port WWN
- SCSI target port WWN
- Flow feature set consisting of Fibre Channel Write Acceleration and statistics monitoring.

SCSI Flow Manager

The SCSI flow manager (SFM) resides on a supervisor module and handles the configuration of SCSI flows, validating them and relaying configuration information to the appropriate SSM. It also handles any dynamic

changes to the status of the SCSI flow due to external events. The SFM registers events resulting from operations, such as port up or down, VSAN suspension, and zoning that affects the SCSI flow status, and updates the flow status and configuration accordingly.

The SFM on the initiator communicates to its peer on the target side using Cisco Fabric Services (CFS). Peer communication allows the initiator SFM to validate target parameters and program information on the target side.

SCSI Flow Configuration Client

A SCSI flow configuration client (SFCC) resides on the CPP of the SSM. It receives flow configuration requests from the SFM, programs the DPP corresponding to the initiator and target port interfaces, and responds to the SFM with the status of the configuration request.

SCSI Flow Data Path Support

The DPP on the SSM examines all the messages between the initiator and target and provides SCSI flow features such as Fibre Channel Write Acceleration and statistics monitoring.



Note

For statistics monitoring, the target device is not required to be connected to an SSM.

Licensing Requirements for SCSI Flow Services

The following table shows the licensing requirements for SCSI Flow Services:

License	License Description
ENTERPRISE_PKG	SCSI flow statistics requires license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you. For a complete explanation of the NX-OS licensing scheme, see the <i>Cisco NX-OS Licensing Guide</i> .
FM_SERVER_PKG	Traffic Analyzer for SCSI flow statistics requires an Enterprise Services license. For a complete explanation of the NX-OS licensing scheme and how to obtain and apply licenses, see the <i>Cisco NX-OS Licensing Guide</i> .

Guidelines and Limitations

The SCSI flow specification is a distributed configuration because the SCSI initiator and the target might be physically connected to SSMs on two different switches located across the fabric. The configuration does not require information to identify either the switch name or the SSM slot location for either the initiator or the target. The manual SCSI flow configuration is performed only at the initiator side. This simplifies the configuration process. The initiator switch sends the configuration to the SFM on the target switch using CFS. No SCSI flow configuration is necessary on the target switch.

Default Settings

[Table 116: Default SCSI Flow Services Parameters](#), on page 952 lists the default settings for SCSI flow services parameters.

Table 116: Default SCSI Flow Services Parameters

Parameters	Default
SCSI flow services	Disabled
SCSI flow services distribution	Enabled

Configuring SCSI Flow Services

This section includes the following topics:

Enabling SCSI Flow Services



Note

Enabling SCSI flow services on interfaces has the following restrictions:

- The fewest number of interfaces that you can enable is four. You can specify fc1 through fc4, but not fc1 through fc2.
- The first interface in the group must be 1, 5, 9, 13, 17, 21, 25, or 29. You can specify fc5 through fc8, but not fc7 through fc10.
- The groups of four interfaces do not need to be consecutive. You can specify fc1 through fc8 and fc17 through fc20.

A SCSI flow identifier is unique on a switch such as VSAN identifiers and is chosen by the user. To configure a SCSI flow identifier, follow these steps:

Procedure

	Command or Action	Purpose
Step 1		

Enabling Intelligent Storage Services



Note

The port range must be a multiple of four (for example fc4/1 through fc4-12).

To configure the values to associate with the fast, medium, and slow migration rates, follow these steps:

Procedure

- Step 1** Expand End Devices and then select SSM Features in the Physical Attributes pane.
You see the Intelligent Storage Services configuration in the Information pane.
- Step 2** Click the SSM tab.
You see the set of configured services in the Information pane.
- Step 3** Click Create Row to enable a new service on an SSM.
You see the Create SSM dialog box.
- Step 4** Select the switch and SSM card you want to configure.
- Step 5** (Optional) Uncheck the Use All Ports on Module check box if you want to provision a subset of the ports on the card to use this service.
- Step 6** Select the port range you want to provision for using this service (starting port and ending port).
- Step 7** Select the feature you want to enable on these ports from the drop-down list of services.
- Step 8** Set the PartnerImageURI field if you are enabling a third-party application that requires an image loaded onto the SSM.
- Step 9** Click Create to create this row and enable this service.
-

Configuring Fibre Channel Using DCNM-SAN

To configure a Fibre Channel using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand End Devices and then select SSM Features in the Physical Attributes pane.
You see the Intelligent Storage Services configuration, showing the FCWA tab in the Information pane.
- Step 2** Click Create Row in the Information pane to create a SCSI flow or click a row in the FCWA table to modify an existing SCSI flow.
You see the FC Write Acceleration dialog box.
- Step 3** Select the initiator and target WWNs and VSAN IDs and check the WriteAcc check box to enable Fibre Channel Write Acceleration on this SCSI flow.
- Step 4** (Optional) Enable SCSI flow statistics on this SCSI flow by checking the Enable Statistics check box.
- Step 5** (Optional) Change the BufCount value to set the number of 2K buffers used by the SCSI target.
- Step 6** Click Create to create this SCSI flow.
-

Disabling Intelligent Storage Services

To disable Intelligent Storage Services in DCNM-SAN for an SSM and free up a group of ports that use these services, follow these steps:

Procedure

- Step 1** Expand End Devices and then select SSM Features in the Physical Attributes pane.
You see the Intelligent Storage Services configuration in the Information pane.
- Step 2** Click the SSM tab.
You see the set of configured services in the Information pane.
- Step 3** Select the row in the table that you want to disable.
- Step 4** (Optional) Check the Reboot Module on Delete check box if you want to force the card to reboot after disabling the service. This is equivalent to the CLI force option.
- Step 5** Click Delete Row. The ports that were provisioned for this service become available for provisioning in another service.

Note If Reboot Module on Delete was checked, then the SSM module reboots.

Verifying SCSI Flow Services

To display SCSI Flow Services configuration information, perform one of the following tasks:

Command	Purpose
show scsi-flow	Displays the SCSI flow services configuration for all specific SCSI flow identifiers.
show ssm provisioning	Displays provisioned applications on an SSM.
show scsi-flow flow-id 3	Displays the SCSI flow services configuration for a specific SCSI flow identifiers.

For detailed information about the fields in the output from these commands, refer to the *Cisco DC-OS Command Reference*.

Displaying SCSI Flow Services Information

Use the **show scsi-flow** command to display information about SCSI flow services.

Displays Applications Provisioned on an SSM

```
switch# show ssm provisioning
Module   Ports      Application      Provisioning Status
-----
4        1-32       scsi-flow       success
```

Displays SCSI Flow Services Configuration for All SCSI Flow Identifiers

```
switch# show scsi-flow
```

```

Flow Id: 3
Initiator VSAN: 101
Initiator WWN: 21:00:00:e0:8b:05:76:28
Target VSAN: 102
Target WWN: 21:00:00:20:37:38:7f:7d
Target LUN: ALL LUNs
Flow Verification Status:
-----
Initiator Verification Status:    success
Target Verification Status:      success
Initiator Linecard Status:       success
Target Linecard Status:         success
Feature Status:
-----
Write-Acceleration enabled
Write-Acceleration Buffers: 1024
Configuration Status:    success
Statistics enabled
Configuration Status:    success

Flow Id: 4
Initiator VSAN: 101
Initiator WWN: 21:00:00:e0:8b:05:76:28
Target VSAN: 102
Target WWN: 21:00:00:20:37:38:a7:89
Target LUN: ALL LUNs
Flow Verification Status:
-----
Initiator Verification Status:    success
Target Verification Status:      success
Initiator Linecard Status:       success
Target Linecard Status:         success
Feature Status:
-----
Write-Acceleration enabled
Write-Acceleration Buffers: 1024
Configuration Status:    success

```

Displays SCSI Flow Services Configuration for a Specific SCSI Flow Identifier

```

switch# show scsi-flow flow-id 3
Flow Id: 3
Initiator VSAN: 101
Initiator WWN: 21:00:00:e0:8b:05:76:28
Target VSAN: 102
Target WWN: 21:00:00:20:37:38:7f:7d
Target LUN: ALL LUNs
Flow Verification Status:
-----
Initiator Verification Status:    success
Target Verification Status:      success
Initiator Linecard Status:       success
Target Linecard Status:         success
Feature Status:
-----
Write-Acceleration enabled
Write-Acceleration Buffers: 1024
Configuration Status:    success
Statistics enabled
Configuration Status:    success

```

Filed Description for SCSI Flow Services

This section includes the following topics:

- SSM
- Virtual Initiator

SSM

Field	Description
StartPort, EndPort, Feature	A table containing feature related information for interfaces. This table gives a list of interfaces that are assigned to different features. The interfaces supported are of the type Fibre Channel.
PartnerImageURI	A collection of objects related to SSM Feature to interface mapping.

Virtual Initiator

Field	Description
Processor Id	The DPP ID.
Control	If false, it's the data path. If true, it's the control path.



CHAPTER 46

Configuring SCSI Flow Statistics

- [Configuring SCSI Flow Statistics](#) , on page 957

Configuring SCSI Flow Statistics

This chapter describes the SCSI flow statistics feature which is supported on the Storage Services Module (SSM).

This chapter includes the following sections:

Information About SCSI Flow Statistics

This section includes the following topics:

SCSI Flow Statistics Overview

The statistics that can be collected for SCSI flows include the following:

- SCSI reads
 - Number of I/Os
 - Number of I/O blocks
 - Maximum I/O blocks
 - Minimum I/O response time
 - Maximum I/O response time
- SCSI writes
 - Number of I/Os
 - Number of I/O blocks
 - Maximum I/O blocks
 - Minimum I/O response time
 - Maximum I/O response time

- Other SCSI commands (not read or write)
 - Test unit ready
 - Report LUN
 - Inquiry
 - Read capacity
 - Mode sense
 - Request sense
- Errors
 - Number of timeouts
 - Number of I/O failures
 - Number of various SCSI status events
 - Number of various SCSI sense key errors or events

To take advantage of this feature, only the initiator must be directly attached to an SSM.



Note The SCSI flow statistics feature requires the Enterprise Package license installed only on the initiator switches.



Note For SCSI flow statistics, the initiator must connect to an SSM on a Cisco MDS switch while the target can connect to any other switch in the fabric. The SCSI flow initiator and target cannot connect to the same switch.

SCSI Flow Specification Attributes

A SCSI flow specification consists of the following attributes:

- SCSI flow identifier
- VSAN identifier
- SCSI initiator port WWN
- SCSI target port WWN
- Flow feature set consisting of Fibre Channel Write Acceleration and statistics monitoring.

SCSI Flow Manager

The SCSI flow manager (SFM) resides on a supervisor module and handles the configuration of SCSI flows, validating them and relaying configuration information to the appropriate SSM. It also handles any dynamic changes to the status of the SCSI flow due to external events. The SFM registers events resulting from operations, such as port up or down, VSAN suspension, and zoning that affects the SCSI flow status, and updates the flow status and configuration accordingly.

The SFM on the initiator communicates to its peer on the target side using Cisco Fabric Services (CFS). Peer communication allows the initiator SFM to validate target parameters and program information on the target side.

SCSI Flow Configuration Client

A SCSI flow configuration client (SFCC) resides on the CPP of the SSM. It receives flow configuration requests from the SFM, programs the DPP corresponding to the initiator and target port interfaces, and responds to the SFM with the status of the configuration request.

SCSI Flow Data Path Support

The DPP on the SSM examines all the messages between the initiator and target and provides SCSI flow features such as Fibre Channel Write Acceleration and statistics monitoring.

Licensing Requirements for SCSI Flow Statistics

The following table shows the licensing requirements for SCSI Flow Statistics:

License	License Requirement
ENTERPRISE_PKG	SCSI flow statistics requires license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you. For a complete explanation of the NX-OS licensing scheme, see the <i>Cisco NX-OS Licensing Guide</i> .
FM_SERVER_PKG	Traffic Analyzer for SCSI flow statistics requires an Enterprise Services license. For a complete explanation of the NX-OS licensing scheme and how to obtain and apply licenses, see the <i>Cisco NX-OS Licensing Guide</i> .

Default Settings

[Table 117: Default SCSI Flow Statistics Parameters](#), on page 959 lists the default settings for SCSI flow statistics parameters.

Table 117: Default SCSI Flow Statistics Parameters

Parameters	Default
SCSI flow statistics	Disabled

Configuring SCSI Flow Statistics

This section includes the following topics:

Enabling SCSI Flow Statistics

To enable SCSI flow statistics monitoring using DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Expand End Devices and then select SSM Features in the Physical Attributes pane.
You see the FCWA tab in the Information pane.
- Step 2** Click Create Row in the Information pane to create a SCSI flow or click a row in the FCWA table to modify an existing SCSI flow.
You see the FC Write Acceleration dialog box.
- Step 3** Select the initiator and target WWNs and VSAN IDs and check the Enable Statistics check box to enable SCSI flow statistics on this SCSI flow.
- Step 4** (Optional) Enable Fibre Channel Write Acceleration on this SCSI flow at this time by checking the WriteAcc check box.
- Step 5** Click Create to create this SCSI flow.
-

Clearing SCSI Flow Statistics

Clears SCSI flow statistics counters for SCSI flow ID.

To clear SCSI flow statistics using DCNM-SAN, follow these steps:

Procedure

-
- Step 1** Expand End Devices and then select SSM Features.
- Step 2** Check the Stats Clear check box to clear SCSI flow statistics.
- Step 3** Click the Apply Changes icon to clear the SCSI flow statistics.
-


Field Descriptions for SCSI Flow Statistics

SSM

Field	Description
StartPort, EndPort, Feature	A table containing feature related information for interfaces. This table gives a list of interfaces that are assigned to different features. The interfaces supported are of the type Fibre Channel.
PartnerImageURI	A collection of objects related to SSM Feature to interface mapping.

Virtual Initiator

Field	Description
Processor Id	The DPP ID.
Control	If false, it's the data path. If true, it's the control path.



CHAPTER 47

Configuring Fibre Channel Write Acceleration

- [Configuring Fibre Channel Write Acceleration, on page 963](#)

Configuring Fibre Channel Write Acceleration

This chapter describes the Fibre Channel Write Acceleration (FC-WA) feature, including how to enable the feature on Cisco NX-OS.

This chapter includes the following sections:

Information About Fibre Channel Write Acceleration

Fibre Channel Write Acceleration minimizes application latency or reduces transactions per second over long distances. For synchronous data replication, Fibre Channel Write Acceleration increases the distance of replication or reduces effective latency to improve performance. With this feature you can also configure the buffer count and change the number of 2-KB buffers reserved on the target side DPP for a SCSI flow.

To take advantage of this feature, both the initiator and target devices must be directly attached to an SSM.

The Fibre Channel Write Acceleration feature also allows the configuration of the buffer count. You can change the number of 2-KB buffers reserved on the target side DPP for a SCSI flow.

You can estimate the number of buffers to configure using the following formula:

$$(\text{Number of concurrent SCSI writes} * \text{size of SCSI writes in bytes}) / \text{FCP data frame size in bytes}$$

For example, HDS TrueCopy between HDS 9970s uses 1-KB FCP data frames. You perform an initial synchronization for a 16-LUN TrueCopy group with 15 tracks, or 768-KB per LUN, which requires approximately $16 * (768 * 1024) / 1024$ or 12248 write buffers.



Note

The Fibre Channel Write Acceleration feature requires the Enterprise Package license installed on both the initiator and target switches.



Note The initiator and target cannot connect to the same Cisco MDS switch. Fibre Channel Write Acceleration requires that the initiator and target must each connect to an SSM module installed on different Cisco MDS switches.



Note Fibre Channel Write Acceleration can only be provisioned on the entire SSM, not a group of interfaces on the SSM.

Licensing Requirements for Fibre Channel Write Acceleration

The following table shows the licensing requirements for Fibre Channel Write Acceleration:

License	License Description
ENTERPRISE_PKG	Fibre Channel Write Acceleration requires license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you. For a complete explanation of the NX-OS licensing scheme, see the <i>Cisco NX-OS Licensing Guide</i> .

Default Settings

[Table 118: Default Fibre Channel Write Acceleration Parameters](#), on page 964 lists the default settings for Fibre Channel Write Acceleration parameters.

Table 118: Default Fibre Channel Write Acceleration Parameters

Parameters	Default
Fibre Channel Write Acceleration	Disabled
Fibre Channel Write Acceleration buffers	1024

Configuring Fibre Channel Write Acceleration

This section includes the following topics:

Enabling Fibre Channel Write Acceleration

To enable Fibre Channel Write Acceleration, and optionally modify the number of Write Acceleration buffers using the DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand End Devices and then select SSM Features from the Physical Attributes pane.

You see the Intelligent Storage Services configuration, showing the FCWA tab in the Information pane.

Step 2 Click Create Row in the Information pane to create a SCSI flow or click a row in the FCWA table to modify an existing SCSI flow.

You see the FC Write Acceleration dialog box.

Step 3 Select the initiator and target WWNs and VSAN IDs and check the WriteAcc check box to enable Fibre Channel Write Acceleration on this SCSI flow.

Step 4 (Optional) Enable SCSI flow statistics on this SCSI flow at this time by checking the Enable Statistics check box.

Step 5 (Optional) Set the BufCount value to the number of 2K buffers used by the SCSI target.

Step 6 Click Create to create this SCSI flow with Fibre Channel Write Acceleration.

Filed Description for Fibre Channel Write Acceleration

This section includes the following topics:

- FCWA
- SSM
- Virtual Initiator
- FCWA Config Status

FCWA

Field	Description
Flow Id	Represents the flow identifier.
Init WWN	Represents the pWWN of the initiator in the flow.
Init VSAN	The VSAN ID of the initiator on which the flow is configured.
Target WWN	Represents the pWWN of the target in the flow.
TargetVSAN	The VSAN ID of the target on which the flow is configured.
WriteAcc	Specifies if write-acceleration feature is enabled for this flow. If set to true it is enabled. If set to false, it is disabled.
BufCount	It specifies the number of buffers to be used for write-acceleration.
Stats Enable	Specifies if the statistics gathering needs to be enabled for this flow. If set to true, then it is enabled. If it is set to false, then it is disabled.
Stats Clear	Assists in clearing the statistics for this flow.
Init Verification	The verification status of the initiator device corresponding to the SCSI flow.
Init Module	The status of the linecard where the SCSI flow initiator device is located.

Field	Description
Target Verification	The verification status of the target device corresponding to the SCSI flow.
Target Module	The status of the linecard where the SCSI flow target device is located.

SSM

Field	Description
StartPort, EndPort, Feature	A table containing feature related information for interfaces. This table gives a list of interfaces that are assigned to different features. The interfaces supported are of the type Fibre Channel.
PartnerImageURI	A collection of objects related to SSM Feature to interface mapping.

Virtual Initiator

Field	Description
Processor Id	The DPP ID.
Control	If false, it's the data path. If true, it's the control path.

FCWA Config Status

Field	Description
Overall	The configuration status for write-acceleration feature for this flow.
Initiator	The initiator configuration status for write-acceleration feature for this flow.
Target	The target configuration status for write-acceleration feature for this flow.



CHAPTER 48

Monitoring the Network

- [Monitoring the Network](#) , on page 967

Monitoring the Network

This chapter describes how the DCNM-SAN manages 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 Network Monitoring

DCNM-SAN provides extensive SAN discovery, topology mapping, and information viewing capabilities. DCNM-SAN collects information on the fabric topology through SNMP queries to the switches connected to it. DCNM-SAN recreates a fabric topology, presents it in a customizable map, and provides inventory and configuration information in multiple viewing options such as fabric view, device view, summary view, and operation view.

Once DCNM-SAN is invoked, a SAN discovery process begins. Using information polled from a seed Cisco MDS 9000 Family switch, including Name Server registrations, Fibre Channel Generic Services (FC-GS), Fabric Shortest Path First (FSPF), and SCSI-3, DCNM-SAN automatically discovers all devices and interconnects on one or more fabrics. All available switches, host bus adapters (HBAs), and storage devices are discovered. The Cisco MDS 9000 Family switches use Fabric-Device Management Interface (FMDI) to retrieve the HBA model, serial number and firmware version, and host operating-system type and version discovery without host agents. DCNM-SAN gathers this information through SNMP queries to each switch. The device information discovered includes device names, software revision levels, vendor, ISLs, PortChannels, and VSANs.

Monitoring Health and Events

DCNM-SAN works with the Cisco MDS 9000 Family switches to show the health and status of the fabric and switches. Information about the fabric and its components is gathered from multiple sources, including Online System Health Management, Call Home, system messages, and SNMP notifications. This information is then made available from multiple menus on DCNM-SAN or Device Manager.

DCNM-SAN Events Tab

The DCM-SAN Events tab, available from the topology window, displays the events DCM-SAN received from sources within the fabric. These sources include SNMP events, RMON events, system messages, and system health messages. The Events tab shows a table of events, including the event name, the source and time of the event, a severity level, and a description of the event. The table is sortable by any of these column headings.

Event Information in DCM-SAN Web Server Reports

The DCM-SAN web server client displays collections of information gathered by the Performance Manager. This information includes events sent to the DCM-SAN 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 DCM-SAN Server. Choose a fabric and then click the **Events** tab to see a summary or detailed report of the events that have occurred in the selected fabric. The summary view shows how many switches, ISLs, hosts, or storage elements are down on the fabric and how many warnings have been logged for that SAN entity. The detailed view shows a list of all events that have been logged from the fabric and can be filtered by severity, time period, or type.

Events in Device Manager

Device Manager displays the events when you choose **Logs > Events**. Device Manager can display the current list of events or an older list of events that has been stored on the DCM-SAN host. The event table shows details on each event, including time, source, severity, and a brief description of the event.

SAN Discovery and Topology Mapping

DCM-SAN provides extensive SAN discovery, topology mapping, and information viewing capabilities. DCM-SAN collects information on the fabric topology through SNMP queries to the switches connected to it. DCM-SAN recreates a fabric topology, presents it in a customizable map, and provides inventory and configuration information in multiple viewing options.

Device Discovery

Once DCM-SAN is invoked, a SAN discovery process begins. Using information polled from a seed Cisco MDS 9000 Family switch, including Name Server registrations, Fibre Channel Generic Services (FC-GS), Fabric Shortest Path First (FSPF), and SCSI-3, DCM-SAN automatically discovers all devices and interconnects on one or more fabrics. All available switches, host bus adapters (HBAs), and storage devices are discovered. The Cisco MDS 9000 Family switches use Fabric-Device Management Interface (FMDI) to retrieve HBA model, serial number and firmware version, and host operating-system type and version discovery without host agents. DCM-SAN gathers this information through SNMP queries to each switch. The device information discovered includes device names, software revision levels, vendor, ISLs, PortChannels, and VSANs.

For a VSAN change involving a third-party switch, DCM-SAN will need a second discovery to show the correct topology due to the discovery dependency when there is any change in a mixed VSAN. The first discovery finds the third-party switch and the subsequent discovery will show the information on which VSAN it is going to join and can discover the end devices connected to it. You can wait for the subsequent discovery or trigger a manual discovery.

Topology Mapping

DCM-SAN is built upon a topology representation of the fabric. DCM-SAN provides an accurate view of multiple fabrics in a single window by displaying topology maps based on device discovery information.

You can modify the topology map icon layout with an easy-to-use, drag-and-drop interface. The topology map visualizes device interconnections, highlights configuration information such as zones, VSANs, and ISLs exceeding utilization thresholds. The topology map also provides a visual context for launching command-line interface (CLI) sessions, configuring PortChannels, and opening device managers.

Using the Topology Map

The DCNM-SAN 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 DCNM-SAN Client for that fabric.

Procedure

-
- | | |
|---------------|--------------------------------------------------------------------------------------------------------------|
| Step 1 | Click File > Preferences to open the DCNM-SAN 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. |
| Step 3 | Click Apply, and then click OK to save this change. |
-

Using Enclosures with DCNM-SAN 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.

Clicking **Alias->Enclosure** displays hosts and storage elements in the Information pane. This is a shortcut to naming enclosures. To use this shortcut, highlight each row in the host or storage table that you want grouped in an enclosure then click Alias -> Enclosure. This automatically sets the enclosure names of each selected row with the first token of the alias.

Mapping Multiple Fabrics

To log into multiple fabrics, the same username and password must be used. The information for both fabrics is displayed, with no need to select a seed switch. To see details of a fabric, click the tab for that fabric at the bottom of the Fabric pane, or double-click the fabric's cloud icon.

Inventory Management

The Information pane in DCNM-SAN shows inventory, configuration, and status information for all switches, links, and hosts in the fabric. Inventory management includes vendor name and model, and software or firmware versions. Select a fabric or VSAN from the Logical Domains pane, and then select the Summary tab in the Information pane to get a count of the number of VSANS, switches, hosts, and storage elements in the fabric.

Viewing Logs from Device Manager

You can view system messages from Device Manager if Device Manager is running from the same workstation as the DCNM-SAN 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.

**Note**

To view syslog local logs, you need to configure the IP address of the DCNM-SAN Server in the syslog host.



CHAPTER 49

Monitoring Performance

- [Monitoring Performance](#) , on page 971

Monitoring Performance

This chapter describes how to configure Performance Monitoring tools for Cisco DCNM-SAN and Device Manager. These tools provide real-time statistics as well as historical performance monitoring.

This chapter contains the following sections:

Information About Performance Monitoring

Device Manager provides an easy tool for monitoring ports on the Cisco MDS 9000 Family switches. This tool gathers statistics at a configurable interval and displays the results in tables or charts. Real-time performance statistics are useful for dynamic troubleshooting and fault isolation within the fabric. Real-time statistics gather data on parts of the fabric in user-configurable intervals and display these results in DCNM-SAN and Device Manager. For a selected port, you can monitor any of a number of statistics including traffic in and out, errors, class 2 traffic, and FICON data.

Real-Time Performance Monitoring

Device Manager provides an easy tool for monitoring ports on the Cisco MDS 9000 Family switches. This tool gathers statistics at a configurable interval and displays the results in tables or charts. These statistics show the performance of the selected port in real-time and can be used for performance monitoring and troubleshooting. For a selected port, you can monitor any of a number of statistics including traffic in and out, errors, class 2 traffic, and FICON data. You can set the polling interval from ten seconds to one hour, and display the results based on a number of selectable options including absolute value, value per second, and minimum or maximum value per second.

Device Manager checking for oversubscription on the host-optimized four-port groups on relevant modules. Right-click the port group on a module and choose Check Oversubscription from the pop-up menu.

Device manager provides two performance views: the Summary View tab and the configurable monitor option per port.

Historical Performance Monitoring

Performance Manager gathers network device statistics historically and provides this information using DCNM-SAN client and 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. See the [“Information About Performance Manager” section on page 6-1](#) for an overview of Performance Manager.

Configuring Performance Manager

This section includes the following topics:

Creating a Flow with Performance Manager

With the Flow Configuration Wizard you can create host-to-storage, storage-to-host, or bidirectional flows. Once defined, you can add these flows to a collection configuration file to monitor the traffic between a host/storage element pair. The flows created become part of the collection options in the Performance Manager Configuration Wizard.

Creating a Collection with Performance Manager

The Performance Manager Configuration Wizard steps you through the process of creating collections using configuration files. Collections are defined for one or all VSANs in the fabric. Collections can include statistics from the SAN element types described in [Table 119: Performance Manager Collection Types, on page 972](#).

Table 119: Performance Manager Collection Types

Collection Type	Description
ISLs	Collects link statistics for ISLs.
Host	Collects link statistics for SAN hosts.
Storage	Collects link statistics for a storage elements.
Flows	Collects flow statistics defined by the Flow Configuration Wizard.

Using Performance Thresholds

The Performance Manager Configuration Wizard allows you to set up two thresholds that trigger events when the monitored traffic exceeds the percent utilization configured. These event triggers can be set as either Critical or Warning events that are reported on the DCNM-SAN web client Events browser page.

You must choose either absolute value thresholds or baseline thresholds that apply to all transmit or receive traffic defined in the collection. Click the **Use absolute values** radio button on the last screen of the Performance Manager Configuration Wizard to configure thresholds that apply directly to the statistics gathered. These statistics, as a percent of the total link capacity, are compared to the percent utilization configured for the threshold type. If the statistics exceed either configured threshold, an event is shown on the DCNM-SAN web client Events tab.

As an example, the collection has absolute value thresholds set for 60% utilization (for warning) and 80% utilization (for critical). If Performance Manager detects that the traffic on a 1-Gigabit link in its collection exceeds 600 Mbps, a warning event is triggered. If the traffic exceeds 800 Mbps, a critical event is triggered.

Baseline thresholds are defined for a configured time of day or week (1 day, 1 week, or 2 weeks). The baseline is created by calculating the average of the statistical results for the configured time each day, week, or every 2 weeks. [Table 120: Baseline Time Periods for a Collection Started on Wednesday at 4pm, on page 973](#) shows an example of the statistics used to create the baseline value for a collection defined at 4 pm on a Wednesday.

Table 120: Baseline Time Periods for a Collection Started on Wednesday at 4pm

Baseline Time Window	Statistics Used in Average Calculation
1 day	Every prior day at 4 pm
1 week	Every prior Wednesday at 4 pm
2 weeks	Every other prior Wednesday at 4 pm

Baseline thresholds create a threshold that adapts to the typical traffic pattern for each link for the same time window each day, week, or every 2 weeks. Baseline thresholds are set as a percent of the average (110% to 500%), where 100% equals the calculated average.

As an example, a collection is created at 4 pm on Wednesday, with baseline thresholds set for 1 week, at 150% of the average (warning) and 200% of the average (critical). Performance Manager recalculates the average for each link at 4 pm every Wednesday by taking the statistics gathered at that time each Wednesday since the collection started. Using this as the new average, Performance Manager compares each received traffic statistic against this value and sends a warning or critical event if the traffic on a link exceeds this average by 150% or 200% respectively.

[Table 121: Example of Events Generated for 1-Gigabit Links, on page 973](#) shows two examples of 1-Gigabit links with different averages in our example collection and at what traffic measurements the Warning and Critical events are sent.

Table 121: Example of Events Generated for 1-Gigabit Links

Average	Warning Event Sent at 150%	Critical Event Sent at 200%
400 Mbps	600 Mbps	800 Mbps
200 Mbps	300 Mbps	400 Mbps

Set these thresholds on the last screen of the Collections Configuration Wizard by checking the **Send events if traffic exceeds threshold** check box.

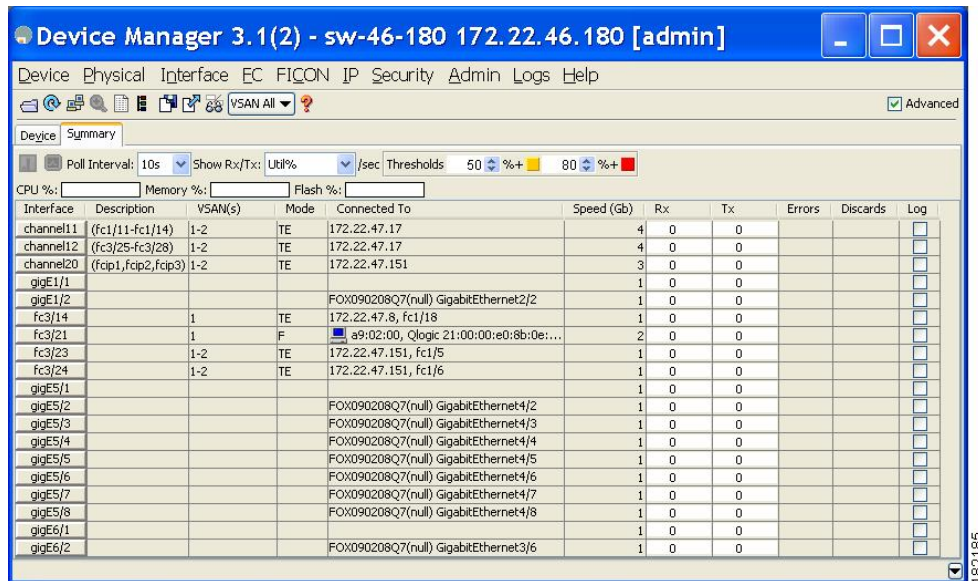
Configuring the Summary View in Device Manager

Procedure

Step 1 Click the Summary tab on the main display.

You see all of the active ports on the switch, as well as the configuration options available from the Summary view .

Figure 164: Device Manager Summary Tab



Step 2 Choose a value from the Poll Interval drop-down list.

Step 3 Decide how you want your data to be interpreted by looking at the Show Rx/Tx drop-down menu. The table updates each polling interval to show an overview of the receive and transmit data for each active port on the switch.

Step 4 Select a value from the **Show Rx/Tx** drop-down list. If you select **Util%**, you need to also select values from the two **Show Rx/Tx > %Util/sec** drop-down lists. The first value is the warning level and the second value is the critical threshold level for event reporting.

Note that you can also display percent utilization for a single port by selecting the port and clicking the Monitor Selected Interface Traffic Util % icon.

Configuring Per Port Monitoring using Device Manager

The configurable monitor per port option gives statistics for in and out traffic on that port, errors, class 2 traffic and other data that can be graphed over a period of time to give a real-time view into the performance of the port.

Procedure

Step 1 Click the **Device** tab.

Step 2 Right-click the port you are interested in and choose Monitor from the drop-down menu.

You see the port real-time monitor dialog box.

Step 3 Select a value from the Interval drop-down list to determine how often data is updated in the table shown here.

Step 4 Click a statistical value in the table then click one of the graphing icons to display a running graph of that statistic over time. You see a graph window that contains options to change the graph type.

Tip You can open multiple graphs for statistics on any of the active ports on the switch.

Displaying DCNM-SAN Real-Time ISL Statistics

This section includes the following topics:

You can configure DCNM-SAN to gather ISL statistics in real time. These ISL statistics include receive and transmit utilization, bytes per second, as well as errors and discards per ISL.

Procedure

Step 1 Choose Performance > ISLs in Real-Time.

You see any ISL statistics in the Information pane.

Figure 165: ISL Performance in Real Time

From Switch	From Interface	To Switch	To Interface	Speed	Rx Util%	Rx Bytes	Rx Pkts	Tx Util%	Tx Bytes	Tx Pkts	Total Errors	Total Discards
sw172-22-46-224	fc1/17	sw172-22-46-221	fc2/17	2 Gb	0	953	7	0	523	9	0	0
sw172-22-46-223	fc1/7	sw172-22-46-222	fc1/7	2 Gb	0	50	0	0	5	0	0	0
sw172-22-46-223	fc1/10	sw172-22-46-222	fc1/10	2 Gb	0	73	1	0	531	5	0	0
sw172-22-46-223	fc1/11	sw172-22-46-222	fc1/11	2 Gb	0	88	1	0	547	5	0	0
sw172-22-46-223	fc1/12	sw172-22-46-222	fc1/12	2 Gb	0	395	6	0	46	1	0	0
sw172-22-46-223	fc1/14	sw172-22-46-222	fc1/14	2 Gb	0	64	0	0	28	0	0	0
sw172-22-46-223	fc1/16	sw172-22-46-221	fc2/16	2 Gb	0	156	2	0	70	1	0	0
sw172-22-46-222	fc1/1	sw172-22-46-221	fc2/29	2 Gb	0	1.308K	20	0	2.148K	17	0	0
sw172-22-46-222	fc1/4	sw172-22-46-221	fc2/4	2 Gb	0	1.026K	13	0	1.648K	16	0	0
sw172-22-46-225	fc1/3	sw172-22-47-118	fc1/20	2 Gb	0	0	0	0	0	0	0	0
sw172-22-46-225	fc1/5	sw172-22-46-224	fc1/5	2 Gb	0	362	3	0	341	4	0	0
sw172-22-46-225	fc1/9	sw172-22-46-224	fc1/9	2 Gb	0	244	3	0	364	4	0	0

Step 2 Select a value from the Poll Interval drop-down list.

Step 3 Select two values from the **Bandwidth** utilization thresholds drop-down lists, one value for the minor threshold and one value for the major threshold.

The table shown updates each polling interval to show the statistics for all configured ISLs in the fabric.

Step 4 Select a row in the table to highlight that ISL in blue in the Topology map.

Viewing Performance Statics Using DCNM-SAN

You can configure DCNM-SAN to gather historic and real time statistics of ISLs or End devices. These statistics include receive and transmit utilization, bytes per second, as well as errors and discards per ISL or end device.

Procedure

Step 1 Right-click the ISL or end device in the Fabric pane.

You see a context menu.

Step 2 Select Show Statics.

Note Show Statics menu will be enabled only if you add the fabric to the Performance Manager collection.

Displaying Performance Manager Reports

This section includes the following topics:

You can view Performance Manager statistical data using preconfigured reports that are built on demand and displayed in a web browser. These reports provide summary information as well as detailed statistics that can be viewed for daily, weekly, monthly, or yearly results.

Procedure

Step 1 Choose Performance > Reports to access Performance Manager reports from DCNM-SAN.

This opens a web browser window showing the default DCNM-SAN web client event summary report.

Step 2 Click the Performance tab to view the Performance Manager reports.

Performance Manager begins reporting data ten minutes after the collection is started.

What to do next



Note DCNM-SAN Web Server must be running for reports to work.

Displaying Performance Summary

The Performance Summary page presents a dashboard display of the throughput and link utilization for hosts, ISLs, storage, and flows for the last 24-hour period. The summary provides a quick overview of the fabric's bandwidth consumption and highlights any hotspots.

The report includes network throughput pie charts and link utilization pie charts. Use the navigation tree on the left to show summary reports for monitored fabrics or VSANs. The summary displays charts for all hosts, storage elements, ISLs, and flows. Each pie chart shows the percent of entities (links, hosts, storage, ISLs, or flows) that measure throughput or link utilization on each of six predefined ranges. Move the mouse over a pie chart section to see how many entities exhibit that range of statistics. Double-click any pie chart to bring up a table of statistics for those hosts, storage elements, ISLs, or flows.

Displaying Performance Tables and Details Graphs

Click Host, Storage, ISL, or Flow to view traffic over the past day for all hosts, storage, ISLs, or flows respectively. A table lists all of the selected entities, showing transmit and receive traffic and errors and discards, if appropriate. The table can be sorted by any column heading. The table can also be filtered by day, week, month, or year. Tables for each category of statistics display average and peak throughput values and provide hot-links to more detailed information.

Clicking a link in any of the tables opens a details page that shows graphs for traffic by day, week, month, and year. If flows exist for that port, you can see which storage ports sent data. The details page also displays graphs for errors and discards if they are part of the statistics gathered and are not zero.

If you double-click a graph on a Detail report, it will launch the Cisco Traffic Analyzer for Fibre Channel, if configured. The aliases associated with hosts, storage devices, and VSANs in the fabric are passed to the Cisco Traffic Analyzer to provide consistent, easy identification.

Displaying Performance of Host-Optimized Port Groups

You can monitor the performance of host-optimized port groups by selecting Performance > End Devices and selecting Port Groups from the Type drop-down list.

Displaying Performance Manager Events

Performance Manager events are viewed through DCNM-SAN Web Server. To view problems and events in DCNM-SAN Web Server, choose a fabric and then click the **Events** tab to see a summary or detailed report of the problems and events that have occurred in the selected fabric.

Generating Performance Manager Reports

Generating Top10 Reports in Performance Manager

You can generate historical Top10 reports that can be saved for later review. These reports list the entities from the data collection, with the most active entities appearing first. This is a static, one-time only report that generates averages and graphs of the data collection as a snapshot at the time the report is generated. These Top10 reports differ from the other monitoring tables and graphs in Performance Manager in that the other data is continuously monitored and is sortable on any table column. The Top10 reports are a snapshot view at the time the report was generated and are static. These are one-time reports that generate averages and graphs of the data collection as a snapshot at the time the report is generated.



Tip Name the reports with a timestamp so that you can easily find the report for a given day or week.

These Top10 reports differ from the other monitoring tables and graphs in Performance Manager in that the other data is continuously monitored and is sortable on any table column. The Top10 reports are a snapshot view at the time the report was generated.



Note Top10 reports require analyzing the existing data over an extended period of time and can take hours or more to generate on large fabrics.

Generating Top10 Reports Using Scripts

You can generate Top10 reports manually by issuing the following commands:

- On UNIX, run the script:

```
"<user_directory>/cisco_mds9000/bin/pm.sh display pm/pm.xml <output_directory>"
```

- On Windows, run the script:

```
"c:\Program Files\Cisco Systems\MDS 9000\bin\pm.bat display pm\pm.xml <output_directory>"
```

On UNIX, you can automate the generation of the Top10 reports on your DCNM-SAN/DCNM-SAN Server host by adding the following cron entry to generate the reports once an hour:

```
0 * * * * /<user_directory>/cisco_mds9000/bin/pm.sh display pm/pm.xml <output_directory>
```

If your crontab does not run automatically or Java complains about an exception similar to the below example, you need to add “-Djava.awt.headless=true” to the JVMARGS command in
/<user_directory>/cisco_mds9000/bin/pm.sh.

Example Java Exception

```
in thread "main" java.lang.InternalError Can't connect to X11 window server using '0.0' as
the value of the DISPLAY variable.
```

Configuring Performance Manager for Use with Cisco Traffic Analyzer

Performance Manager works in conjunction with the Cisco Traffic Analyzer to allow you to monitor and manage the traffic on your fabric. Using Cisco Traffic Analyzer with Performance Manager requires the following components:

- A configured Fibre Channel Switched Port Analyzer (SPAN) destination (SD) port to forward Fibre Channel traffic.
- A Port Analyzer Adapter 2 (PAA-2) to convert the Fibre Channel traffic to Ethernet traffic.
- Cisco Traffic Analyzer software to analyze the traffic from the PAA-2.

Procedure

-
- Step 1** Set up the Cisco Traffic Analyzer according to the instructions in the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*.
- Step 2** Get the following three items of information:
- The IP address of the management workstation on which you are running Performance Manager and Cisco Traffic Analyzer.
 - The path to the directory where Cisco Traffic Analyzer is installed.
 - The port that is used by Cisco Traffic Analyzer (the default is 3000).
- Step 3** Start the Cisco Traffic Analyzer.
- a) Choose **Performance > Traffic Analyzer > Open**.
 - b) Enter the URL for the Cisco Traffic Analyzer, in the format:

Example:

```
http://<ip address>
>:<port number>
>
```

ip address is the address of the management workstation on which you have installed the Cisco Traffic Analyzer

:port number is the port that is used by Cisco Traffic Analyzer (the default is :3000).

- c) Click **OK**.
- d) Choose **Performance > Traffic Analyzer > Start**.
- e) Enter the location of the Cisco Traffic Analyzer, in the format:

Example:

```
D:\<directory>
>\ntop.bat
```

D: is the drive letter for the disk drive where the Cisco Traffic Analyzer is installed.

directory is the directory containing the ntop.bat file.

- f) Click **OK**.

Step 4 Create the flows you want Performance Manager to monitor, using the Flow Configuration Wizard. See the [Creating a Flow with Performance Manager, on page 972](#)

Step 5 Define the data collection you want Performance Manager to gather, using the Performance Manager Configuration Wizard. See the [Creating a Collection with Performance Manager, on page 972](#).

- a) Choose the VSAN you want to collect information for or choose All VSANs.
- b) Check the types of items you want to collect information for (Hosts, ISLs, Storage Devices, and Flows).
- c) Enter the URL for the Cisco Traffic Analyzer in the format:

Example:

```
http://<ip address>/<directory>
```

where:

ip address is the address of the management workstation on which you have installed the Cisco Traffic Analyzer, and *directory* is the path to the directory where the Cisco Traffic Analyzer is installed.

- d) Click **Next**.
- e) Review the data collection on this and the next section to make sure this is the data you want to collect.
- f) Click **Finish** to begin collecting data.

Note Data is not collected for JBOD or for virtual ports. If you change the data collection configuration parameters during a data collection, you must stop and restart the collection process for your changes to take effect.

Step 6 Choose **Performance > Reports** to generate a report. Performance Manager Web Server must be running. You see Web Services; click **Custom** then select a report template.

Note It takes at least five minutes to start collecting data for a report. Do not attempt to generate a report in Performance Manager during the first five minutes of collection.

Step 7 Click **Cisco Traffic Analyzer** at the top of the Host or Storage detail pages to view the Cisco Traffic Analyzer information, or choose **Performance > Traffic Analyzer > Open**. The Cisco Traffic Analyzer page will not open unless ntop has been started already.

- Note** For information on capturing a SPAN session and starting a Cisco Traffic Analyzer session to view it, refer to the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*.
- Note** For information on viewing and interpreting your Performance Manager data, see the [Creating a Flow with Performance Manager, on page 972](#). For information on viewing and interpreting your Cisco Traffic Analyzer data, refer to the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*.

What to do next

For performance drill-down, DCNM-SAN Server can launch the Cisco Traffic Analyzer in-context from the Performance Manager graphs. The aliases associated with hosts, storage devices, and VSANs are passed to the Cisco Traffic Analyzer to provide consistent, easy identification.

Exporting Data Collections

This section includes the following topics:

Exporting Data Collections to XML Files

The RRD files used by Performance Manager can be exported to a freeware tool called rrdtool. The rrd files are located in pm/db on the DCNM-SAN Server. To export the collection to an XML file, enter the following command at the operating system command-line prompt:

```
/bin/pm.bat xport xxx yyy
```

In this command, xxx is the RRD file and yyy is the XML file that is generated. This XML file is in a format that rrdtool is capable of reading with the command:

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

You can import an XML file with the command:

```
bin/pm.bat pm restore <xmlFile> <rrdFile>
```

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

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

The pm xport and pm restore commands can be found on your DCNM-SAN Server at bin\PM.bat for Windows platforms or bin/PM.sh on UNIX platforms. For more information on the rrdtool, refer to the following website: <http://www.rrdtool.org>.

Exporting Data Collections in Readable Format

You can export the RRD files used by Performance Manager to a freeware tool called rrdtool and export the collection to an XML file. Cisco MDS SAN-OS Release 2.1(1a) introduces the inability to export data collections in comma-separated format (CSV). This format can be imported to various tools, including Microsoft Excel. You can export these readable data collections either from the DCNM-SAN Web Services

menus or in batch mode from the command line on Windows or UNIX. Using DCNM-SAN Web Services, you can export one file. Using batch mode, you can export all collections in the pm.xml file.



Note DCNM-SAN Web Server must be running for this to work.

You can export data collections to Microsoft Excel using DCNM-SAN Web Server.

Procedure

-
- Step 1** Click the Performance tab on the main page.
You see the overview table.
- Step 2** Click the **Flows** sub-tab.
- Step 3** Right-click the name of the entity you want to export and select **Export to Microsoft Excel**.
You see the Excel chart for that entity in a pop-up window.
-

Exporting Data Collections in Readable Format

You can export data collections using command-line batch mode.

Procedure

-
- Step 1** Go to the installation directory on your workstation and then go to the bin directory.
- Step 2** On Windows, enter `.\pm.bat export C:\Program Files\Cisco Systems\MDS 9000\pm\pm.xml <export directory>`. This creates the csv file (export.csv) in the export directory on your workstation.
- Step 3** On UNIX, enter `./pm.sh export /usr/local/cisco_mds9000/pm/pm.xml <export directory>`. This creates the csv file (export.csv) in the export directory on your workstation.
-

What to do next

When you open this exported file in Microsoft Excel, the following information displays:

- Title of the entity you exported and the address of the switch the information came from.
- The maximum speed seen on the link to or from this entity.
- The VSAN ID and maximum speed.
- The timestamp, followed by the receive and transmit data rates in bytes per second.

Analyzing SAN Health

The SAN Health Advisor tool is a utility that used to monitor the performance and collect the statistics. You can perform the following tasks with this tool:

- Run Performance Monitor to collect I/O statistics
- Collect fabric inventory (switches and other devices)
- Create a graphical layout of fabric topology
- Create reports of error conditions and statistical data

You can install this tool at any SAN environment to collect I/O statistics for the specified time (usually 24 hours), generate health reports and automatically send reports to the designated system administrator for review at regular intervals.

When you start SAN Health Advisor tool, it runs in wizard mode, and prompts for inputs such as seed switch credentials, IP address of the server to which the data to be sent and all the necessary information for the software setup. As soon as the fabric is discovered, the tool starts capturing performance data, I/O statistics and error conditions.

The reports generated from the collection is stored in the \$INSTALLDIR/dcm/fm/reports directory. These reports are automatically sent to the designated SAN administrator for review. In a situation where the tool fails to collect the data, it generates a report with an error message or exception. After sending the reports the tool automatically uninstalls itself and terminates all the processes that it established on the host machine.

The report that SAN Health Advisor tool generates will have the following details:

- Events
- System messages
- Analysis of connectivity
- Zone discrepancy
- System configuration
- Interface status
- Domain information
- Security settings

Installing the SAN Health Advisor Tool

SAN Health Advisor tool can be installed and run on Windows, UNIX, and Solaris platforms. Install the package that contains the .jar file with JRE version 6.0.



Note The SAN Health tool is not installed by default when you install DCNM-SAN software.

Procedure

- Step 1** Double-click the San Health Advisor tool installer.
You see the San Health Advisor tool Installer window.
- Step 2** Select an installation folder on your workstation for SAN Health Advisor.
On Windows, the default location is C:\Program Files\Cisco Systems\.
- Step 3** Click Install to start the installation.
You will see the installation progressing.
You will see the Fabric Options dialog box

- Step 4** In the Seed Switch text box, enter the IP address of the seed switch.
- Step 5** Enter the user name and password for the switch.
- Step 6** Select the authentication privacy option from the Auth-Privacy drop-down list box.
- Step 7** Click the Performance Collection check box to enable the process to run for 24 hours.
- Step 8** Click Collect to start gathering performance information.
- You see the collecting dialog box.
- If you want to stop gathering information in the middle of the process, click Cancel. You see the message indicating performance collection is complete.
- Step 9** Click Uninstall to remove the SAN Health Advisor software.
-



CHAPTER 50

Configuring Call Home

- [Configuring Call Home, on page 985](#)

Configuring Call Home

Call Home provides e-mail-based notification of critical system events. A versatile range of message formats are available for optimal compatibility with pager services, standard e-mail, or XML-based automated parsing applications.



Note

Cisco Autonotify is upgraded to a new capability called Smart Call Home. Smart Call Home has significant functionality improvement over Autonotify and is available across the Cisco product range. For detailed information on Smart Call Home, see the Smart Call Home page at this location:
<http://www.cisco.com/go/smartcall/>

This chapter includes the following sections:

Information About Call Home

The Call Home feature provides message throttling capabilities. Periodic inventory messages, port syslog messages, and RMON alert messages are added to the list of deliverable Call Home messages. If required you can also use the Cisco Fabric Services application to distribute the Call Home configuration to all other switches in the fabric.

The Call Home service provides e-mail-based notification of critical system events. A versatile range of message formats are available for optimal compatibility with pager services, standard e-mail, or XML-based automated parsing applications.

Common features may include the following:

- Paging the network support engineer
- E-mailing the Network Operations Center
- Raising a direct case with the Technical Assistance Center

The Call Home functionality is available directly through the Cisco MDS 9000 Family switches and the Cisco Nexus 5000 Series switches. It provides multiple Call Home messages, each with separate potential destinations.

You can define your own destination profiles in addition to predefined profiles; you can configure up to 50 e-mail addresses for each destination profile. Flexible message delivery and format options make it easy to integrate specific support requirements.

The Call Home feature offers the following advantages:

- Fixed set of predefined alerts for trigger events on the switch.
- Automatic execution and attachment of relevant command output.

This section includes the following topics:

Call Home Features

The Call Home functionality is available directly through the Cisco MDS 9000 Family switches and the Cisco Nexus 5000 Series switches. It provides multiple Call Home profiles (also referred to as *Call Home destination profiles*), each with separate potential destinations. You can define your own destination profiles in addition to predefined profiles.

The Call Home function can even leverage support from Cisco Systems or another support partner. Flexible message delivery and format options make it easy to integrate specific support requirements.

The Call Home feature offers the following advantages:

- Fixed set of predefined alerts and trigger events on the switch.
- Automatic execution and attachment of relevant command output.
- Multiple message format options:
 - Short Text—Suitable for pagers or printed reports.
 - Plain Text—Full formatted message information suitable for human reading.
 - XML—Matching readable format using Extensible Markup Language (XML) and document type definitions (DTDs) named Messaging Markup Language (MML). The MML DTD is published on the Cisco.com website at <http://www.cisco.com/> . The XML format enables communication with the Cisco Systems Technical Assistance Center.
- Multiple concurrent message destinations. You can configure up to 50 e-mail destination addresses for each destination profile.
- Multiple message categories including system, environment, switching module hardware, supervisor module, hardware, inventory, syslog, RMON, and test.
- Secure messages transport directly from your device or through an HTTP proxy server or a downloadable transport gateway (TG). You can use a TG aggregation point to support multiple devices, or in cases where security requires that your devices not be connected directly to the Internet.

About Smart Call Home

Smart Call Home is a component of Cisco SMARTnet Service that offers proactive diagnostics, real-time alerts, and personalized web-based reports on select Cisco devices.

Smart Call Home provides fast resolution of system problems by analyzing Call Home messages sent from your devices and providing a direct notification path to Cisco customer support.

Smart Call Home offers the following features:

- Continuous device health monitoring and real-time diagnostics alerts.
- Analysis of Call Home messages from your device and where appropriate, automatic service request generation, routed to the appropriate TAC team, including detailed diagnostic information to speed problem resolution.
- Web-based access to Call Home messages and recommendations, inventory and configuration information for all Call Home devices. Provides access to associated Field Notices, Security Advisories and End-of-Life Information.

[Table 122: Benefits of Smart Call Home Compared to Autonotify](#) , on page 987 lists the benefits of Smart Call Home.

Table 122: Benefits of Smart Call Home Compared to Autonotify

Feature	Smart Call Home	Autonotify
Low touch registration	The registration process is considerably streamlined. Customers no longer need to know their device serial number or contract information. They can register devices without manual intervention from Cisco by sending a message from those devices. The procedures are outlined at www.cisco.com/go/smartcall	Requires the customer to request Cisco to add each specific serial number to the database.
Recommendations	Smart Call Home provides recommendations for known issues including those for which SRs are raised and for which SRs are not appropriate but for which customers might want to still take action on.	Autonotify raises SRs for a set of failure scenarios but no recommendations are provided for these.
Device report	Device report includes full inventory and configuration details. Once available, the information in these reports will be mapped to field notices, PSIRTs, EoX notices, configuration best practices and bugs.	No.
History report	The history report is available to look up any message and its contents, including show commands, message processing, analysis results, recommendations and service request numbers for all messages sent over the past three months.	A basic version is available that does not include contents of message.
Network summary report	A report that provides a summary of the make-up of devices and modules in the customer network (for those devices registered with Smart Call home)	No.
Cisco device support	Device Support will be extended across the Cisco product range. See the supported products table at www.cisco.com/go/smartcall	Deprecated in favor of Smart Call Home in October 2008.

Obtaining Smart Call Home

If you have a service contract directly with Cisco Systems, you can receive automatic case generation from the Technical Assistance Center by registering with the Smart Call Home service.

You need the following items to register:

- The SMARTnet contract number for your switch.
- Your e-mail address
- Your Cisco.com ID

For detailed information on Smart Call Home, including quick start configuration and registration steps, see the Smart Call Home page at this location:

<http://www.cisco.com/go/smartcall/>

Call Home Destination Profiles

A destination profile contains the required delivery information for an alert notification. Destination profiles are typically configured by the network administrator.

Using alert groups you can select the set of Call Home alerts to be received by a destination profile (predefined or user defined). Alert groups are predefined subsets of Call Home alerts supported in all switches in the Cisco MDS 9000 Family and the Cisco Nexus 5000 Series. Different types of Call Home alerts are grouped into different alert groups depending on their type. You can associate one or more alert groups to each profile as required by your network.

Call Home Alert Groups

An alert group is a predefined subset of Call Home alerts supported in all switches in the Cisco MDS 9000 Family and Cisco Nexus 5000 Series. Alert groups allow you to select the set of Call Home alerts to be received by a destination profile (predefined or user-defined). A Call Home alert is sent to e-mail destinations in a destination profile only if that Call Home alert belongs to one of the alert groups associated with that destination profile.

Using the predefined Call Home alert groups you can generate notification messages when certain events occur on the switch. You can customize predefined alert groups to execute additional **show** commands when specific events occur and to notify you of output other than from the predefined **show** commands.

Customized Alert Group Messages

An alert group is a predefined subset of Call Home alerts supported in all switches in the Cisco MDS 9000 Family and Cisco Nexus 5000 Series switches. Alert groups allow you to select the set of Call Home alerts to be received by a destination profile (predefined or user-defined). The predefined Call Home alert groups generate notification messages when certain events occur on the switch. You can customize predefined alert groups to execute additional show commands when specific events occur.

The output from these additional **show** commands is included in the notification message along with the output of the predefined **show** commands.

Call Home Message Level Feature

The Call Home message level feature allows you to filter messages based on their level of urgency. Each destination profile (predefined and user-defined) is associated with a Call Home message level threshold. Any message with a value lower than the urgency threshold is not sent. Call Home severity levels are not the same as system message logging severity levels.

Syslog-Based Alerts

You can configure the switch to send certain syslog messages as Call Home messages. The messages are sent based on the mapping between the destination profile and the alert group mapping, and on the severity level of the generated syslog message.

To receive a syslog-based Call Home alert, you must associate a destination profile with the syslog alert groups (currently there is only one syslog alert group—syslog-group-port) and configure the appropriate message level.

The syslog-group-port alert group selects syslog messages for the port facility. The Call Home application maps the syslog severity level to the corresponding Call Home severity level (see the [Call Home Message Levels, on page 993](#)). For example, if you select level 5 for the Call Home message level, syslog messages at levels 0, 1, and 2 are included in the Call Home log.

Whenever a syslog message is generated, the Call Home application sends a Call Home message depending on the mapping between the destination profile and the alert group mapping and based on the severity level of the generated syslog message. To receive a syslog-based Call Home alert, you must associate a destination profile with the syslog alert groups (currently there is only one syslog alert group—syslog-group-port) and configure the appropriate message level (see the [Call Home Message Levels, on page 993](#)).



Note Call Home does not change the syslog message level in the message text. The syslog message texts in the Call Home log appear as they are described in the Cisco MDS 9000 Family and Nexus 7000 Series System Messages Reference.

RMON-Based Alerts

You can configure the switch to send Call Home notifications corresponding to RMON alert triggers. All RMON-based Call Home messages have their message level set to NOTIFY (2). The RMON alert group is defined for all RMON-based Call Home alerts. To receive an RMON-based Call Home alert, you must associate a destination profile with the RMON alert group.

General E-Mail Options Using HTTPS Support

The HTTPS support for Call Home provides a transport method called HTTP. HTTPS support is used for a secure communication, and HTTP is used for nonsecure communication. You can configure an HTTP URL for the Call Home destination profile as a destination. The URL link can be from a secure server or nonsecure server. For a destination profile configured with the HTTP URL, the Call Home message is posted to the HTTP URL link.



Note The Call Home HTTP configuration can be distributed over CFS on the switches running NX-OS Release 4.2(1) and later. The Call Home HTTP configuration cannot be distributed to switches that support the nondistributable HTTP configuration. Switches running lower versions than NX-OS Release 4.2(1) and later will ignore the HTTP configuration.

Multiple SMTP Server Support

Cisco MDS NX-OS and Cisco NX-OS 5000 Series switches support multiple SMTP servers for Call Home. Each SMTP server has a priority configured between 1 and 100, with 1 being the highest priority and 100 being the lowest. If the priority is not specified, a default value of 50 is used.

You can configure up to five SMTP servers for Call Home. The servers are contacted based on their priority. The highest priority server is contacted first. If the message fails to be sent, the next server in the list is contacted until the limit is exhausted. If two servers have equal priority, the one that was configured earlier is contacted.

If a high-priority SMTP server fails, the other servers will be contacted. A time delay may occur while sending a message. The delay is minimal if the attempt to send the message through the first SMTP server is successful. The delay may increase depending on the number of unsuccessful attempts with different SMTP servers.

**Note**

The new configuration process is not related to the old configuration. However, if the SMTP servers are configured using both the old and new schemes, the older configuration is of the highest priority.

Multiple SMTP servers can be configured on any MDS 9000 Family switch, Cisco Nexus 5000 Series switches, and Cisco Nexus 7000 Series switches running Release 5.0(1a) or later.

The new configuration will only be distributed to switches that have multiple SMTP servers. The older switches in the fabric will ignore the new configuration received over CFS.

In a mixed fabric that has CFS enabled, the switches running NX-OS Release 5.0 can configure new functionalities and distribute the new configuration to other switches with Release 5.0 in the fabric over CFS. However, if an existing switch running NX-OS Release 4.x upgrades to Release 5.0, the new configurations will not be distributed to that switch as a CFS merge is not triggered on an upgrade. There are two options to upgrade:

- Apply new configuration only when all the switches in the fabric support them. (Recommended option).
- Do an empty commit from an existing NX-OS Release 5.0 switch which has the new configuration

Periodic Inventory Notification

You can configure the switch to periodically send a message with an inventory of all software services currently enabled and running on the switch along with hardware inventory information. The inventory is modified each time the switch is restarted nondisruptively.

Duplicate Message Throttle

You can configure a throttling mechanism to limit the number of Call Home messages received for the same event. If the same message is sent multiple times from the switch within a short period of time, you may be swamped with a large number of duplicate messages.

Call Home Configuration Distribution

You can enable fabric distribution for all Cisco MDS 9000 Family switches and Cisco Nexus 5000 Series switches in the fabric. When you perform Call Home configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric. However, the switch priority and the Syscontact names are not distributed.

You automatically acquire a fabric-wide lock when you enter the first configuration command operation after you enable distribution in a switch. The Call Home application uses the effective and pending database model to store or commit the configuration changes. When you commit the configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. After making the configuration changes, you can choose to discard the changes by aborting the changes instead of committing them. In either case, the lock is released. See [Chapter 13, “Using the CFS Infrastructure .”](#) for more information on the CFS application.

**Note**

The switch priority and the Syscontact name are not distributed.

Fabric Lock Override

If you have performed a Call Home task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.

**Tip**

The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

Clearing Call Home Name Server Database

When the Call Home name server database is full, a new entry cannot be added. The device is not allowed to come online. To clear the name server database, increase the database size or perform a cleanup by removing unused devices. A total of 20,000 name server entries are supported.

EMC E-mail Home Delayed Traps

DCNM-SAN can be configured to generate EMC E-mail Home XML e-mail messages. In SAN-OS Release 3.x or earlier, DCMN-SAN listens to interface traps and generates EMC E-mail Home e-mail messages. Link traps are generated when an interface goes to down from up or vice versa. For example, if there is a scheduled server reboot, the link goes down and DCMN-SAN generates an e-mail notification.

Cisco NX-OS Release 4.1(3) provides the ability to generate a delayed trap so that the number of generated e-mail messages is reduced. This method filters server reboots and avoids generating unnecessary EMC E-mail Home e-mail messages. In NX-OS Release 4.1(3), users have the ability to select the current existing feature or this new delayed trap feature.

Event Triggers

This section discusses Call Home trigger events. Trigger events are divided into categories, with each category assigned CLI commands to execute when the event occurs. The command output is included in the transmitted message. [Table 123: Event Triggers , on page 992](#) lists the trigger events.

Table 123: Event Triggers

Event	Alert Group	Event Name	Description	Call Home Message Level
Call Home	System and CISCO_TAC	SW_CRASH	A software process has crashed with a stateless restart, indicating an interruption of a service.	5
	System and CISCO_TAC	SW_SYSTEM_INCONSISTENT	Inconsistency detected in software or file system.	5
	Environmental and CISCO_TAC	TEMPERATURE_ALARM	Thermal sensor indicates temperature reached operating threshold.	6
		POWER_SUPPLY_FAILURE	Power supply failed.	6
		FAN_FAILURE	Cooling fan has failed.	5
	Line Card Hardware and CISCO_TAC	LINECARD_FAILURE	Line card hardware operation failed.	7
		POWER_UP_DIAGNOSTICS_FAILURE	Line card hardware failed power-up diagnostics.	7
	Line Card Hardware and CISCO_TAC	PORT_FAILURE	Hardware failure of interface port(s).	6
	Line Card Hardware, Supervisor Hardware, and CISCO_TAC	BOOTFLASH_FAILURE	Failure of boot compact flash card.	6
	Supervisor Hardware and CISCO_TAC	NVRAM_FAILURE	Hardware failure of NVRAM on supervisor hardware.	6
	Supervisor Hardware and CISCO_TAC	FREEDISK_FAILURE	Free disk space is below a threshold on supervisor hardware.	6
	Supervisor Hardware and CISCO_TAC	SUP_FAILURE	Supervisor hardware operation failed.	7
		POWER_UP_DIAGNOSTICS_FAILURE	Supervisor hardware failed power-up diagnostics.	7
	Supervisor Hardware and CISCO_TAC	INBAND_FAILURE	Failure of in-band communications path.	7

Event	Alert Group	Event Name	Description	Call Home Message Level
	Supervisor Hardware and CISCO_TAC	EOBC_FAILURE	Ethernet out-of-band channel communications failure.	6
Call Home	Supervisor Hardware and CISCO_TAC	MGMT_PORT_FAILURE	Hardware failure of management Ethernet port.	5
	License	LICENSE_VIOLATION	Feature in use is not licensed, and are turned off after grace period expiration.	6
Inventory	Inventory and CISCO_TAC	COLD_BOOT	Switch is powered up and reset to a cold boot sequence.	2
		HARDWARE_INSERTION	New piece of hardware inserted into the chassis.	2
		HARDWARE_REMOVAL	Hardware removed from the chassis.	2
Test	Test and CISCO_TAC	TEST	User generated test.	2
Port syslog	Syslog-group-port	SYSLOG_ALERT	Syslog messages corresponding to the port facility.	2
RMON	RMON	RMON_ALERT	RMON alert trigger messages.	2

Call Home Message Levels

Table 124: Event Categories and Executed Commands

Event Category	Description	Executed Commands
System show module show version show tech-support platform show tech-support sysmgr show hardware show sprom all	Events generated by failure of a software system that is critical to unit operation.	show tech-support show system redundancy status

Event Category	Description	Executed Commands
Environmental show module show version show environment show logging logfile tail -n 200	Events related to power, fan, and environment sensing elements such as temperature alarms.	show moduleshow environment
Line Card Hardware show module show version show tech-support platform show tech-support sysmgr show hardware show sprom all	Events related to standard or intelligent line card hardware.	show tech-support
Supervisor Hardware show module show version show tech-support platform show tech-support sysmgr show hardware show sprom all	Events related to supervisor modules.	show tech-support
Inventory show module show version show hardware show inventory show system uptime show sprom all show license usage	Inventory status is provided whenever a unit is cold booted, or when FRUs are inserted or removed. This is considered a noncritical event, and the information is used for status and entitlement.	show version
Test show module show version	User generated test message.	show version

Call Home messages (sent for syslog alert groups) have the syslog severity level mapped to the Call Home message level (see the [Syslog-Based Alerts, on page 989](#)).

This section discusses the severity levels for a Call Home message when using one or more switches in the Cisco MDS 9000 Family and the Cisco Nexus 5000 Series. Call Home message levels are preassigned per event type.

Severity levels range from 0 to 9, with 9 having the highest urgency. Each syslog level has keywords and a corresponding syslog level.



Note Call Home does not change the syslog message level in the message text. The syslog message texts in the Call Home log appear as they are described in the Cisco MDS 9000 Family and Nexus 7000 Series System Messages Reference.



Note Call Home severity levels are not the same as system message logging severity levels (see the Cisco MDS 9000 Family and Nexus 7000 Series System Messages Reference).

Table 125: Severity and Syslog Level Mapping

Call Home Level	Keyword Used	Syslog Level	Description
Catastrophic (9)	Catastrophic	N/A	Network wide catastrophic failure.
Disaster (8)	Disaster	N/A	Significant network impact.
Fatal (7)	Fatal	Emergency (0)	System is unusable.
Critical (6)	Critical	Alert (1)	Critical conditions, immediate attention needed.
Major (5)	Major	Critical (2)	Major conditions.
Minor (4)	Minor	Error (3)	Minor conditions.
Warning (3)	Warning	Warning (4)	Warning conditions.
Notify (2)	Notification	Notice (5)	Basic notification and informational messages. Possibly independently insignificant.
Normal (1)	Normal	Information (6)	Normal event signifying return to normal state.
Debug (0)	Debugging	Debug (7)	Debugging messages.

Message Contents

The following contact information can be configured on the switch:

- Name of the contact person
- Phone number of the contact person
- E-mail address of the contact person

- Mailing address to which replacement parts must be shipped, if required
- Site ID of the network where the site is deployed
- Contract ID to identify the service contract of the customer with the service provider

[Table 126: Short Text Messages](#), on page 996 describes the short text formatting option for all message types.

Table 126: Short Text Messages

Data Item	Description
Device identification	Configured device name
Date/time stamp	Time stamp of the triggering event
Error isolation message	Plain English description of triggering event
Alarm urgency level	Error level such as that applied to system message

[Table 127: Reactive Event Message Format](#), on page 996, [Table 128: Inventory Event Message Format](#), on page 999, and [Table 129: User-Generated Test Message Format](#), on page 1001 display the information contained in plain text and XML messages.

Table 127: Reactive Event Message Format

Data Item(Plain text and XML)	Description(Plain text and XML)	XML Tag (XML only)
Time stamp	Date and time stamp of event in ISO time notation: <i>YYYY-MM-DDTHH:MM:SS</i> . Note The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time - ch:EventType
Message name	Name of message. Specific event names are listed in the Event Triggers , on page 991.	/mml/header/name
Message type	Specifically “Call Home.”	/mml/header/type - ch:Type
Message group	Specifically “reactive.”	/mml/header/group
Severity level	Severity level of messag.	/mml/header/level - aml-block:Severity
Source ID	Product type for routing.	/mml/header/source - ch:Series

Data Item(Plain text and XML)	Description(Plain text and XML)	XML Tag (XML only)
Device ID	<p>Unique device identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format is <i>type@Sid@serial</i>, where:</p> <ul style="list-style-type: none"> • <i>type</i> is the product model number from backplane SEEPROM. • @ is a separator character. • <i>Sid</i> is “C,” identifying the serial ID as a chassis serial number. • <i>serial</i> is the number identified by the Sid field. <p>Example: DS-C9509@C@12345678</p>	/mml/ header/deviceId
Customer ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/header/customerID - ch:CustomerId
Contract ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/header/contractId - ch:ContractId>
Site ID	Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/header/siterId - ch:SiteId
Server ID	<p>If the message is generated from the fabric switch, it is the unique device identifier (UDI) of the switch.</p> <p>Format is <i>type@Sid@serial</i>, where:</p> <ul style="list-style-type: none"> • <i>type</i> is the product model number from backplane SEEPROM. • @ is a separator character. • <i>Sid</i> is “C,” identifying the serial ID as a chassis serial number. • <i>serial</i> is the number identified by the Sid field. <p>Example: DS-C9509@C@12345678</p>	/mml/header/serverId - -blank-
Message description	Short text describing the error.	/mml/body/msgDesc - ch:MessageDescription
Device name	Node that experienced the event. This is the host name of the device.	/mml/body/sysName - ch:SystemInfo/Name
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact - ch:SystemInfo/Contact
Contact e-mail	E-mail address of person identified as contact for this unit.	/mml/body/sysContacte-mail - ch:SystemInfo/Contacte-mail

Data Item(Plain text and XML)	Description(Plain text and XML)	XML Tag (XML only)
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhoneNumber - ch:SystemInfo/ContactPhoneNumber
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress - ch:SystemInfo/StreetAddress
Model name	Model name of the switch. This is the specific model as part of a product family name.	/mml/body/chassis/name - rme:Chassis/Model
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo - rme:Chassis/SerialNumber
Chassis part number	Top assembly number of the chassis.	/mml/body/fru/partNo - rme:chassis/Card/PartNumber
Chassis hardware version	Hardware version of chassis.	/mml/body/chassis/hwVersion - rme:Chassis/HardwareVersion
Supervisor module software version	Top level software version.	/mml/body/fru/swVersion - rme:chassis/Card/SoftwareIdentity
Affected FRU name	Name of the affected FRU generating the event message.	/mml/body/fru/name - rme:chassis/Card/Model
Affected FRU serial number	Serial number of affected FRU.	/mml/body/fru/serialNo - rme:chassis/Card/SerialNumber
Affected FRU part number	Part number of affected FRU.	/mml/body/fru/partNo - rme:chassis/Card/PartNumber
FRU slot	Slot number of FRU generating the event message.	/mml/body/fru/slot - rme:chassis/Card/LocationWithinContainer
FRU hardware version	Hardware version of affected FRU.	/mml/body/fru/hwVersion - rme:chassis/Card/SoftwareIdentity
FRU software version	Software version(s) running on affected FRU.	/mml/body/fru/swVersion - rme:chassis/Card/SoftwareIdentity
Command output name	The exact name of the issued command.	/mml/attachments/attachment/name - aml-block:Attachment/Name
Attachment type	Specifically command output.	/mml/attachments/attachment/type - aml-block:Attachment type
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment/mime - aml-block:Attachment/Data encoding
Command output text	Output of command automatically executed.	/mml/attachments/attachment/atdata - aml-block:Attachment/Data

Table 128: Inventory Event Message Format

Data Item(Plain text and XML)	Description(Plain text and XML)	XML Tag(XML only)
Time stamp	Date and time stamp of event in ISO time notation: <i>YYYY-MM-DDTHH:MM:SS</i> . Note The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time - ch:EventTime
Message name	Name of message. Specifically “Inventory Update” Specific event names are listed in the Event Triggers, on page 991 .	/mml/header/name
Message type	Specifically “Inventory Update.”	/mml/header/type - ch-inv:Type
Message group	Specifically “proactive.”	/mml/header/group
Severity level	Severity level of inventory event is level 2.	/mml/header/level - aml-block:Severity
Source ID	Product type for routing at Cisco. Specifically “MDS 9000.”	/mml/header/source - ch-inv:Series
Device ID	Unique Device Identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format is <i>type@Sid@serial</i> , where: <ul style="list-style-type: none">• <i>type</i> is the product model number from backplane SEEPROM.• @ is a separator character.• <i>Sid</i> is “C,,” identifying the serial ID as a chassis serial number.• <i>serial</i> is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/ header /deviceId
Customer ID	Optional user-configurable field used for contact info or other ID by any support service.	/mml/header/customerID - ch-inv:CustomerId
Contract ID	Optional user-configurable field used for contact info or other ID by any support service.	/mml/header/contractId - ch-inv:ContractId>
Site ID	Optional user-configurable field, can be used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/header/siteId - ch-inv:SiteId

Data Item(Plain text and XML)	Description(Plain text and XML)	XML Tag(XML only)
Server ID	<p>If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch.</p> <p>Format is <i>type@Sid@serial</i>, where:</p> <ul style="list-style-type: none"> <i>type</i> is the product model number from backplane SEEPROM. @ is a separator character. <i>Sid</i> is “C,,” identifying the serial ID as a chassis serial number. <i>serial</i> is the number identified by the Sid field. <p>Example: DS-C9509@C@12345678</p>	/mml/header/serverId - -blank-
Message description	Short text describing the error.	/mml/body/msgDesc - ch-inv:MessageDescription
Device name	Node that experienced the event.	/mml/body/sysName - ch-inv:SystemInfo/Name
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact - ch-inv:SystemInfo/Contact
Contact e-mail	E-mail address of person identified as contact for this unit.	/mml/body/sysContacte-mail - ch-inv:SystemInfo/Contacte-mail
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhoneNumber - ch-inv:SystemInfo/ContactPhoneNumber
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress - ch-inv:SystemInfo/StreetAddress
Model name	Model name of the unit. This is the specific model as part of a product family name.	/mml/body/chassis/name - rme:Chassis/Model
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo - rme:Chassis/SerialNumber
Chassis part number	Top assembly number of the chassis.	/mml/body/fru/partNo - rme:chassis/Card/PartNumber
Chassis hardware version	Hardware version of chassis.	/mml/body/fru/hwVersion - rme:chassis/Card/SoftwareIdentity
Supervisor module software version	Top level software version.	/mml/body/fru/swVersion - rme:chassis/Card/SoftwareIdentity
FRU name	Name of the affected FRU generating the event message.	/mml/body/fru/name - rme:chassis/Card/Model

Data Item(Plain text and XML)	Description(Plain text and XML)	XML Tag(XML only)
FRU s/n	Serial number of FRU.	/mml/body/fru/serialNo - rme:chassis/Card/SerialNumber
FRU part number	Part number of FRU.	/mml/body/fru/partNo - rme:chassis/Card/PartNumber
FRU slot	Slot number of FRU.	/mml/body/fru/slot - rme:chassis/Card/LocationWithinContainer
FRU hardware version	Hardware version of FRU.	/mml/body/fru/hwVersion - rme:chassis/Card/SoftwareIdentity
FRU software version	Software version(s) running on FRU.	/mml/body/fru/swVersion - rme:chassis/Card/SoftwareIdentity
Command output name	The exact name of the issued command.	/mml/attachments/attachment/name - aml-block:Attachment/Name
Attachment type	Specifically command output.	/mml/attachments/attachment/type - aml-block:Attachment type
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment/mime - aml-block:Attachment/Data encoding
Command output text	Output of command automatically executed after event categories (see Event Triggers, on page 991).	/mml/attachments/attachment/atdata - aml-block:Attachment/Data

Table 129: User-Generated Test Message Format

Data Item(Plain text and XML)	Description(Plain text and XML)	XML Tag(XML only)
Time stamp	Date and time stamp of event in ISO time notation: <i>YYYY-MM-DDTHH:MM:SS</i> . Note The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time - ch:EventTime
Message name	Name of message. Specifically test message for test type message. Specific event names listed in the Event Triggers, on page 991).	/mml/header/name
Message type	Specifically “Test Call Home.”	/mml/header/type - ch:Type
Message group	This field should be ignored by the receiving Call Home processing application, but may be populated with either “proactive” or “reactive.”	/mml/header/group
Severity level	Severity level of message, test Call Home message.	/mml/header/level - aml-block:Severity

Data Item(Plain text and XML)	Description(Plain text and XML)	XML Tag(XML only)
Source ID	Product type for routing.	/mml/header/source - ch:Series
Device ID	<p>Unique device identifier (UDI) for end device generating message. This field should empty if the message is nonspecific to a fabric switch. Format is <i>type@Sid@serial</i>, where:</p> <ul style="list-style-type: none"> <i>type</i> is the product model number from backplane SEEPROM. <i>@</i> is a separator character. <i>Sid</i> is “C” identifying the serial ID as a chassis serial number. <i>serial</i> is the number identified by the Sid field. <p>Example: DS-C9509@C@12345678</p>	/mml/ header /deviceId
Customer ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/header/customerID - ch:CustomerId
Contract ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/header/contractId - ch:ContractId
Site ID	Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/header/siteId - ch:SiteId
Server ID	<p>If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch.</p> <p>Format is <i>type@Sid@serial</i>, where:</p> <ul style="list-style-type: none"> <i>type</i> is the product model number from backplane SEEPROM. <i>@</i> is a separator character. <i>Sid</i> is “C” identifying the serial ID as a chassis serial number. <i>serial</i> is the number identified by the Sid field. <p>Example: “DS-C9509@C@12345678</p>	/mml/header/serverId - -blank-
Message description	Short text describing the error.	/mml/body/msgDesc - ch:MessageDescription
Device name	Switch that experienced the event.	/mml/body/sysName - ch:SystemInfo/Name
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact - ch:SystemInfo/Contact

Data Item(Plain text and XML)	Description(Plain text and XML)	XML Tag(XML only)
Contact e-mail	E-mail address of person identified as contact for this unit.	/mml/body/sysContacte-mail - ch:SystemInfo/Contacte-mail
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhoneNumber - ch:SystemInfo/ContactPhoneNumber
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress - ch:SystemInfo/StreetAddress
Model name	Model name of the switch. This is the specific model as part of a product family name.	/mml/body/chassis/name - rme:Chassis/Model
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo - rme:Chassis/SerialNumber
Chassis part number	Top assembly number of the chassis. For example, 800-xxx-xxxx.	/mml/body/fru/partNo - rme:chassis/Card/PartNumber
Command output text	Output of command automatically executed after event categories.	/mml/attachments/attachment/atdata - aml-block:Attachment/Data
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment/mime - aml-block:Attachment/Data encoding
Attachment type	Specifically command output.	/mml/attachments/attachment/type - aml-block:Attachment type
Command output name	The exact name of the issued command.	/mml/attachments/attachment/name - aml-block:Attachment/Name

Guidelines and Limitations

Call Home Database Merger Guidelines

When merging two Call Home databases, follow these guidelines:

- Be aware that the merged database contains the following information:
 - A superset of all the destination profiles from the dominant and subordinate switches that take part in the merge protocol.
 - The e-mail addresses and alert groups for the destination profiles.
 - Other configuration information (for example, message throttling, periodic inventory) from the switch that existed in the dominant switch before the merge.
- Verify that two destination profiles do not have the same name (even if they have different configuration information) on the subordinate and dominant switches. If they do contain the same name, the merge operation will fail. You must then modify or delete the conflicting destination profile on the required switch.

See the [“CFS Merge Support” section on page 13-5](#) for detailed concepts.

Call Home Configuration Guidelines

When configuring Call Home, follow these guidelines:

- An e-mail server and at least one destination profile (predefined or user-defined) must be configured. The destination profile(s) used depends on whether the receiving entity is a pager, e-mail, or automated service such as Cisco Smart Call Home.
- Switches can forward events (SNMP traps/informs) up to 10 destinations.
- The contact name (SNMP server contact), phone, and street address information must be configured before Call Home is enabled. This configuration is required to determine the origin of messages received.
- The Cisco MDS 9000 Family switch and the Cisco Nexus 5000 Series switch must have IP connectivity to an e-mail server.
- If Cisco Smart Call Home is used, an active service contract must cover the device being configured.

Default Settings

[Table 130: Default Call Home Default Settings](#), on page 1004 lists the default Call Home settings.

Table 130: Default Call Home Default Settings

Parameters	Default
Destination message size for a message sent in full text format.	500,000
Destination message size for a message sent in XML format.	500,000
Destination message size for a message sent in short text format.	4000
DNS or IP address of the SMTP server to reach the server if no port is specified.	25
Alert group association with profile.	All
Format type.	XML
Call Home message level.	0 (zero)
HTTP proxy server use.	Disabled and no proxy server configured.
HTTP proxy server message size for full text destination.	1 MB
HTTP proxy server message size for XML.	1 MB

Configuring Call Home

How you configure the Call Home process depends on how you intend to use the feature.

This section includes the following topics:

Task Flow for Configuring Call Home

Follow these steps to configure Call Home:

Procedure

- Step 1** Configure contact information.
 - Step 2** Enable or disable Call Home.
 - Step 3** Configure destination profiles.
 - Step 4** Associate one or more alert groups to each profile as required by your network. Customize the alert groups, if desired.
 - Step 5** Configure e-mail options.
 - Step 6** Test Call Home messages.
-

Configuring Contact Information

Switch priority is configured by a user for each switch in the fabric. This priority is used by the operations personnel or TAC support personnel to decide which Call Home message they should respond to first. You can prioritize Call Home alerts of the same severity from each switch.

To assign the contact information, follow these steps:

Before you begin

Each switch must include e-mail, phone, and street address information. You can optionally include the contract ID, customer ID, site ID, and switch priority information.

Procedure

- Step 1** Expand Events and select Call Home from the Physical Attributes pane.
You see the Call Home tabs in the Information pane.
- Step 2** In Device Manager, click Admin > Events > Call Home.
- Step 3** Click the **General** tab, then assign contact information and enable the Call Home feature. Call Home is not enabled by default. You must enter an e-mail address that identifies the source of Call Home notifications.
- Step 4** Click the **Destination(s)** tab to configure the destination e-mail addresses for Call Home notifications. You can identify one or more e-mail addresses that will receive Call Home notifications.

Note Switches can forward events (SNMP traps/informs) up to 10 destinations.

- a) Click the Create tab to create a new destination. You will see the create destination window.
- b) Enter the profile name, ID, and type of destination. You can select email or http in the Type field.

If you select email, you can enter the e-mail address in the EmailAddress field. The HttpUrl field is disabled.

If you select http, you can enter the HTTP URL in the HttpUrl field. The EmailAddress field is disabled.

c) Click Create to complete the destination profile creation.

Step 5 Click the **e-mail Setup** tab to identify the SMTP server. Identify a message server to which your switch has access. This message server will forward the Call Home notifications to the destinations.

Step 6 In DCNM-SAN, click the **Apply Changes** icon. In Device Manager, click **Apply**.

Enabling Call Home Function

Once you have configured the contact information, you must enable the Call Home function.

To enable the Call Home function, follow these steps:

Procedure

Step 1 Select a switch in the Fabric pane.

Step 2 Expand Events and select Call Home in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Click the **Control** tab.

Step 4 Select a switch in the information pane.

Step 5 Check the Duplicate Message Throttle check box.

Step 6 Click the Apply Changes icon.

Configuring Destination Profiles

A destination profile contains the required delivery information for an alert notification. Destination profiles are typically configured by the network administrator.

You can configure the following attributes for a destination profile:

- Profile name—A string that uniquely identifies each user-defined destination profile and is limited to 32 alphanumeric characters. The format options for a user-defined destination profile are full-txt, short-txt, or XML (default).
- Destination address—The actual address, pertinent to the transport mechanism, to which the alert should be sent.
- Message formatting—The message format used for sending the alert (full text, short text, or XML).



Note

If you use the Cisco Smart Call Home service, the XML destination profile is required (see http://www.cisco.com/en/US/partner/products/hw/ps4159/ps4358/products_configuration_example09186a0080108e72.shtml).

Configuring Destination Profiles Messaging options

To configure predefined destination profile messaging options, follow these steps:

Before you begin

At least one destination profile is required. You can configure multiple destination profiles of one or more types. You can use one of the predefined destination profiles or define a desired profile. If you define a new profile, you must assign a profile name.

Procedure

-
- | | |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Expand Events and select Call Home in the Physical Attributes pane. |
| Note | The Destination tab is disabled until you click the Profiles tab. The profiles have to be loaded for the destination tab to be populated. |
| Step 2 | Click the Profiles tab in the Information pane.
You see the Call Home profiles for multiple switches. |
| Step 3 | Set the profile name, message format, message size, and severity level. |
| Step 4 | Click in the Alert Groups column and select or remove an alert group. |
| Step 5 | Click the Apply Changes icon to create this profile on the selected switches. |
-

Configuring Destination Profiles related parameters

To configure a new destination-profile (and related parameters), follow these steps:

Procedure

-
- | | |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Expand Events and select Call Home in the Physical Attributes pane. |
| Note | The Destination tab is disabled until you click the Profiles tab. The profiles have to be loaded for the destination tab to be populated. |
| Step 2 | Click the Profiles tab in the Information pane. You see Call Home profiles for multiple switches. |
| Step 3 | Click the Create Row icon to add a new profile. |
| Step 4 | Set the profile name, message format, size, and severity level. |
| Step 5 | Click an alert group and select each group that you want sent in this profile. |
| Step 6 | Click a transport method. You can select email, http or email and http. |
| Step 7 | Click Create to create this profile on the selected switches. |
-

Associating an Alert Group

Different types of Call Home alerts are grouped into different alert groups depending on their type. You can associate one or more alert groups to each profile as required by your network.

The alert group feature allows you to select the set of Call Home alerts to be received by a destination profile (either predefined or user-defined). You can associate multiple alert groups with a destination profile.

Restrictions

- A Call Home alert is sent to e-mail destinations in a destination profile only if that Call Home alert belongs to one of the alert groups associated with that destination profile.

To associate an alert group with a destination profile, follow these steps:

Procedure

-
- Step 1** Expand Events and select Call Home in the Physical Attributes pane.
 - Step 2** Click the Profiles tab in the Information pane. You see the Call Home profiles for multiple switches.
 - Step 3** Click the **Alert Groups** column in the row for the profile you want to associate. You see the alert groups drop-down menu.
 - Step 4** Click an alert group to select it for association.
 - Step 5** You see a check next to that alert group. To deselect it and remove the check, click it again.
 - Step 6** Click the **Apply Changes** icon.
-

Customizing Alert Group Messages

To assign **show** commands to be executed when an alert is sent, you must associate the commands with the alert group. When an alert is sent, Call Home associates the alert group with an alert type and attaches the output of the **show** commands to the alert message.



Note Make sure the destination profiles for a non-Cisco-TAC alert group, with a predefined **show** command, and the Cisco-TAC alert group are not the same.

Restrictions

- You can assign a maximum of five user-defined **show** commands to an alert group. Only **show** commands can be assigned to an alert group.
- Customized **show** commands are only supported for full text and XML alert groups. Short text alert groups (short-txt-destination) do not support customized **show** commands because they only allow 128 bytes of text.

To customize Call Home alert group messages, follow these steps:

Procedure

-
- Step 1** Expand Events and select Call Home in the Physical Attributes pane.
 - Step 2** Click the User Defined Command tab in the Information pane. You see the User Defined Command information.
 - Step 3** Click the **Create Row** icon.
 - Step 4** Check the check boxes in front of the switches from which you want to receive alerts.
 - Step 5** Select the alert grouptype from the Alert Group Type drop-down list.
 - Step 6** Select the ID (1-5) of the CLI command. The ID is used to keep track of the messages.
 - Step 7** Enter the CLI **show** command in the **CLI Command** field.

- Step 8** Click **Create**.
 - Step 9** Repeat Steps 3 through 7 for each command you want to associate with the profile.
 - Step 10** Click **Close** to close the dialog box.
-

Setting the Call Home Message Levels

The urgency level ranges from 0 (lowest level of urgency) to 9 (highest level of urgency), and the default is 0 (all messages are sent).

To set the message level for each destination profile for Call Home, follow these steps:

Procedure

- Step 1** Expand Events and select Call Home in the Physical Attributes pane. You see the Call Home information in the Information pane. In Device Manager, choose Admin > Events > Call Home.
 - Step 2** Click the Profiles tab in the Information Pane. You see the Call Home profiles.
 - Step 3** Set a message level for each switch using the drop-down menu in the MsgLevel column.
 - Step 4** Click the Apply Changes icon to save your changes.
-

Configuring the Syslog-Based Alerts

To configure the syslog-group-port alert group, follow these steps:

Procedure

- Step 1** Select a switch in the Fabric pane.
 - Step 2** Expand Events and select Call Home in the Physical Attributes pane. You see the Call Home information in the Information pane.
 - Step 3** Click the **Profiles** tab. You see the Call Home profiles.
 - Step 4** Click the **Create Row** icon. You see the Create Call Home Profile dialog box.
 - Step 5** Select the switches for which you want to send alerts.
 - Step 6** Enter the name of the profile in the Name field.
 - Step 7** Choose the message format, message size, and message severity level.
 - Step 8** Check the **syslogGroupPort** check box in the AlertGroups section.
 - Step 9** Click **Create** to create the profile for the syslog-based alerts.
 - Step 10** Close the dialog box.
-

Configuring RMON Alerts

To configure RMON alert groups, follow these steps:

Procedure

- Step 1** Select a switch in the Fabric pane.
 - Step 2** Expand Events and select Call Home in the Physical Attributes pane. You see the Call Home information in the Information pane.
 - Step 3** Click the **Profiles** tab. You see the Call Home profiles.
 - Step 4** Select the **Create Row** icon. You see the Create Call Home Profile dialog box.
 - Step 5** Select switches to send alerts.
 - Step 6** Enter the name of the profile.
 - Step 7** Select the message format, message size, and message severity level.
 - Step 8** Check the **RMON** check box in the Alert Groups section.
 - Step 9** Click **Create** to create the profile for the RMON-based alerts.
 - Step 10** Close the dialog box.
-

Configuring General E-Mail Options

You can configure the from, reply-to, and return-receipt e-mail addresses. While most e-mail address configurations are optional, you must configure the SMTP server address for the Call Home functionality to work.

To configure general e-mail options, follow these steps:

Procedure

- Step 1** Select a switch in the Fabric pane.
 - Step 2** Expand Events and select Call Home in the Physical Attributes pane. You see the Call Home information in the Information pane.
 - Step 3** Click the **e-mail Setup** tab.
 - Step 4** Select a switch in the Information pane.
 - Step 5** Enter the general e-mail information.
 - Step 6** Enter the SMTP server IP address type, IP address or name, and port.
 - Step 7** Click the **Apply Changes** icon to update the e-mail options.
-

Configuring HTTPS Support

Any predefined or user-defined destination profiles can be configured with the HTTPS URL address.

Enable or Disable Transport Method

Any predefined or user-defined destination profiles can be configured to enable or disable a particular transport method. The transport methods are HTTP and e-mail.

Configuring an HTTP Proxy Server

Beginning with Cisco NX-OS Release 5.2, you can configure Smart Call Home to send HTTP messages through an HTTP proxy server. If you do not configure an HTTP proxy server, Smart Call Home sends HTTP messages directly to the Cisco Transport Gateway (TG).



Note The default value for full text destination and for XML is 1 MB.

To configure a Call Home HTTP proxy server, follow these steps:

Procedure

- Step 1** Select a switch in the Fabric pane.
 - Step 2** Expand Events, select Call Home, and HTTP Proxy Server in the Physical Attributes pane. You see the Call Home HTTP Proxy Server information in the Information pane.
 - Step 3** Click the Address Typetab. The Address Type options are displayed.
 - Step 4** Click the Address tab and enter the address of the HTTP proxy server.
 - Step 5** Click the Port tab and enter a integer number to specify the port of the HTTP proxy server.
 - Step 6** Check the Enable check box to enable the HTTP proxy configured for Call Home.
 - Step 7** (Optional) Set an empty value in the Address tab to delete the HTTP proxy server from the MDS switch.
 - Step 8** Choose an address type. You can select ipv4, ipv6, or DNS.
Note If the address is empty, then no proxy server is configured.
 - Step 9** Click **Apply** to update HTTP Proxy Server options.
-

Configuring Call Home Wizard

Task Flow for Configuring Call Home Wizard

Follow these steps to configure the Call Home Wizard:

Procedure

- Step 1** Configure contact information.
 - Step 2** Configure SMTP information.
 - Step 3** Configure the email source and destination information.
 - Step 4** Use CFS to populate the configuration data.
 - Step 5** Display the status.
-

Launching Call Home Wizard

Before You Begin

- Enable the global CFS on the switch from DCNM-SAN configuration table.
- Clear the CFS lock on the switch.
- Check the merger status of CFS on the switch. If a merger failure is found, the wizard clears up the merge failure in the backend process while running the wizard.

To configure Call Home wizard, follow these steps:

Procedure

-
- | | |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Select a fabric in the logical domain tree. |
| Step 2 | Select Tools, Events and Call Home. The master switch pane is displayed. |
| Step 3 | (Optional) You can also launch the Call Home wizard by clicking the CallHome Wizard icon in the Call Home Control tab. |
| Step 4 | Select a Master Switch and click Next. The contact information pane is displayed. |
| Step 5 | Enter the Contact, Phone Number, Email Address and the Street Address information. |
| | Note You must specify all of the four parameters before clicking Next. |
| Step 6 | Click Next. The Email Setup pane is displayed. |
| Step 7 | In the Email SMTP Servers tab, enter the Primary SMTP Server address. You can specify up to two SMTP servers if the master switch is version 5.0 or above. However, you cannot specify a secondary SMTP server if the master switch version is below 5.0. The wizard creates new rows in the SMTP server table. |
| Step 8 | In the Destination tab, click Add to enter the Call Home destinations. You can enter up to three Call Home destinations. |
| Step 9 | (Optional) Click Remove to delete a Call Home destination entry. |
| Step 10 | From the drop-down list, select Protocol and Profile. The Profile drop-down lists three default profiles: xml, short_txt and full_txt. |
| Step 11 | Click Finish to configure the wizard. The Status Dialog window is displayed. All major configuration steps and failures are displayed in the Status Dialog window. |
| Step 12 | Click Run Test to perform the Call Home test. |
| Step 13 | Click Yes to test the command on all switches in the selected fabric or click No to close the window. |
-

Enabling Periodic Inventory Notifications

When you enable this feature without configuring an interval value, the Call Home message is sent every 7 days. This value ranges from 1 to 30 days. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family and Cisco Nexus 5000 Series switches.

To enable periodic inventory notification in a Cisco MDS 9000 Family switch or a Cisco Nexus 5000 Series switch, follow these steps:

Procedure

- Step 1** Select a switch in the Fabric pane.
 - Step 2** Expand Events and select Call Home in the Physical Attributes pane. You see the Call Home information in the Information pane.
 - Step 3** Click the **Periodic Inventory** tab. You see the Call Home periodic inventory information.
 - Step 4** Select a switch in the Information pane.
 - Step 5** Check the **Enable** check box.
 - Step 6** Enter the number of days for which you want the inventory checked.
 - Step 7** Click the **Apply Changes** icon.
-

Configuring Duplicate Message Throttle

You can configure a throttling mechanism to limit the number of Call Home messages received for the same event. If the same message is sent multiple times from the switch within a short period of time, you may be swamped with a large number of duplicate messages.

Restrictions

- By default, this feature is enabled in all switches in the Cisco MDS 9000 Family and the Cisco Nexus 5000 Series switches. When enabled, if the number of messages sent exceeds the maximum limit of 30 messages within the 2-hour time frame, then additional messages for that alert type are discarded within that time frame. You cannot modify the time frame or the message counter limit.
- If 2 hours have elapsed since the first such message was sent and a new message has to be sent, then the new message is sent and the time frame is reset to the time when the new message was sent and the count is reset to 1.

To enable message throttling in a Cisco MDS 9000 Family switch or a Cisco Nexus 5000 Series switch, follow these steps:

Procedure

- Step 1** Select a switch in the Fabric pane.
 - Step 2** Expand Events and select Call Home in the Physical Attributes pane. You see the Call Home information in the Information pane.
 - Step 3** Click the **Control** tab.
 - Step 4** Select a switch in the Information pane.
 - Step 5** Check the **Duplicate Msg Throttle** check box.
 - Step 6** Click the **Apply Changes** icon.
-

Enabling Call Home Fabric Distribution

To enable Call Home fabric distribution, follow these steps:

Procedure

-
- Step 1** Select a switch in the Fabric pane.
 - Step 2** Expand Events and select Call Home in the Physical Attributes pane. You see the Call Home information in the Information pane.
 - Step 3** Click the **CFS** tab. You see the CFS information for Call Home.
 - Step 4** Select a switch in the Information pane.
 - Step 5** Select **Enable** from the drop-down list in the Admin column in the row for that switch.
 - Step 6** Click the **Apply Changes** icon to commit the changes.
-

Call Home Communications Test

You can test Call Home communications by sending a test message to the configured destination(s) or sending a test inventory message to the configured destination(s).

Use the **test** command to simulate a message generation.

To test the Call Home function and simulate a message generation, follow these steps:

Procedure

-
- Step 1** Select a switch in the Fabric pane.
 - Step 2** Expand Events and select Call Home in the Physical Attributes pane. You see the Call Home information in the Information pane.
 - Step 3** Click the **Test** tab. You see the configured tests for the switch and the status of the last testing.
 - Step 4** Select a switch in the Information pane.
 - Step 5** From the TestAction drop-down list in the row for that switch, select **test** or **testWithInventory**
 - Step 6** Click the **Apply Changes** icon to run the test.
-

What to do next

[Table 131: EMC Call Home Traps, on page 1014](#) includes all the traps for EMC Call Home.

Table 131: EMC Call Home Traps

SNMP Trap	Send EMC Call Home When
connUnitStatusChange	operStatus == failed(5)
cefcModuleStatusChange	operStatus != {ok(2), boot(5), selfTest(6), poweredUp(16), syncInProgress(21)}
cefcPowerStatusChange	operStatus = {offDenied(4), offEnvPower(5), offEnvTemp(6), offEnvFan(7), failed(8)}
cefcFRURemoved	all

SNMP Trap	Send EMC Call Home When
cefcFanTrayStatusChange	all
cieDelayedLinkUpDown	operStatusReason != {linkFailure, adminDown, portGracefulShutdown}
cefcFRUInserted	all
entSensorThresholdNotification	value >= threshold

Configuring Delayed Traps

The `server.callhome.delayedtrap.enable` property is added to section 9 Call Home in the `server.properties` configuration file. The property file can enable the DCNM-SAN server to use delayed traps instead of regular linkDown traps for EMC E-mail Home messages.

Enable delayed traps

To enable delayed traps on switches running NX-OS Release 4.1(3) and later, follow these steps:

Before you begin

To enable this feature, you need to turn on delayed traps at switch level, and then set the `server.callhome.delayedtrap.enable` property in the `server.properties` configuration file to true. By default, the `server.callhome.delayedtrap.enable` option is disabled and regular linkDown traps are used.

Procedure

-
- Step 1** Expand Events and select SNMP Traps in the Physical Attributes pane. In the table above the map layout in DCNM-SAN, click the Delayed Traps tab.
 - Step 2** Check the Enable check box for the switches on which you want to enable delayed traps.
 - Step 3** Enter the timer value in the Delay column.
 - Step 4** Click Apply to save your changes.

Note If no value is entered, the default value of 4 minutes is used.

Disable delayed traps

To disable delayed traps, follow these steps:

Procedure

-
- Step 1** Uncheck the Enable check box.
 - Step 2** Click Apply.
-

Enabling Delayed Traps Using Cisco Device Manager

To enable the delayed trap feature, follow these steps:

Procedure

-
- Step 1** In Device Manager, choose Admin > Events > Filters > Delayed Traps. You can see the Events Filters information in the Information pane.
- Step 2** Click the Delayed Traps tab.
- Step 3** Check the Enable check box to enable delayed traps. Delay interval will only be available when the feature is enabled.
- Step 4** To disable Delayed Traps, uncheck the Enable check box and click Apply.
-

Viewing Event Filter Notification

To see the descriptive notification, follow these steps:

In Device Manager, choose Admin > Events > Filters. You can see the Event Filters information in the Information pane. The Event Filters screen displays the descriptive information about the notification.

Field Descriptions for Call Home

This section describes the field descriptions for Call Home.

Call Home General

Field	Description
Contact	The contact person for this switch, together with information on how to contact this person.
PhoneNumber	The phone number of the contact person. The phone number must start with '+' and contains only numeric characters except for space and '-'. Some valid phone numbers are +44 20 8332 9091 +45 44886556 +81-46-215-4678 +1-650-327-2600.
EmailAddress	The e-mail address of the contact person. Some valid e-mail addresses are raj@helpme.com, bob@service.com, mtom@abc.caview.ca.us.
StreetAddress	The mailing address of this switch.
CustomerId	A string, in whatever format is appropriate, to identify the customer.
ContractId	A string, in whatever format is appropriate, to identify the support contract between the customer and support partner.
SiteId	A location identifier of this device.
DeviceServicePriority	The service priority of the device. This determines how fast the device has to be serviced.

Field	Description
Enable	Enables/disables the Call Home infrastructure on the local device.

Related Topics

[Information About Call Home, on page 985](#)

Call Home Destinations

Field	Description
E-mailAddress	The e-mail address associated this destination profile. Some examples are raj@helpme.com, bob@service.com, mtom@abc.caview.ca.us.

Related Topics

[Call Home Destination Profiles, on page 988](#)

Call Home SMTP Servers

Field	Description
Address Type, Address	IP address of the SMTP server.
Port	TCP port of the SMTP server.
Priority	Priority value.

Call Home E-mail Setup

Field	Description
From	The e-mail address that is to be used in the From field when sending the e-mail using SMTP. Some examples are raj@helpme.com, bob@service.com, mtom@abc.caview.ca.us.
ReplyTo	The e-mail address that is to be used in the Reply-To field when sending the e-mail using SMTP. Some examples are raj@helpme.com, bob@service.com, mtom@abc.caview.ca.us.
IP Address Type	The IP address type (IPv4, IPv6, or DNS).
Name or IP Address	Name or IP address of the SMTP server.
Port	TCP port of the SMTP server.

Related Topics

[Configuring General E-Mail Options, on page 1010](#)

Call Home Alerts

Field	Description
Action	Test — Sends a Call Home message TestWithInventory — Sends a message with inventory details.
Status	The status of the last Call Home action invocation.
FailureCause	The failure cause for the last Call Home test invocation.
LastTimeSent	When the last Call Home alert was sent.
NumberSent	The number of Call Home alerts sent.
Interval	Time frame for sending the periodic software inventory Call Home message.
Throttling Enable	If checked, enables the message throttling mechanism implemented on the system, to limit the number of Call Home messages for an alert type within a time frame. The maximum is 30 in a 2-hour time frame, and any further messages for that alert type are discarded.
Enable	If checked, enables the sending of periodic software inventory Call Home messages on the system.

Related Topics

[Call Home Alert Groups, on page 988](#)

[Customizing Alert Group Messages, on page 1008](#)

Call Home User Defined Command

Field	Description
User Defined Command	Configures user-defined commands for the Call Home alert group types.

Delayed Traps

Field	Description
Enable	Enables or disables delay traps.
Delay	Delays interval in minutes (valid values are between 1 to 60).

Call Home Profiles

Field	Description
MsgFormat	XML, full text, or short text.
MaxMsgSize	Maximum message size that can be sent to destination pointed to by this destination profile.

Field	Description
MsgLevel	Threshold level, used for filtering alert messages sent to a destination. Callhome alert message with severity level lower than the configured threshold level would not be sent. The default threshold level is debug (1), which means all the alert messages will be sent.
AlertGroups	The list of configured alert groups for this destination profile.

Event Destinations Addresses

Field	Description
Address/Port	IP address and port to send event.
Security Name	The SNMP parameters to be used when generating messages to be sent to this address.
Security Model	Is used when generating SNMP messages using this entry.
Inform Type	<ul style="list-style-type: none"> • Trap — Unacknowledged event • Inform — Acknowledged event.
Inform Timeout	This expected maximum round-trip time for communicating with the address.
RetryCount	The number of retries to be attempted when a response is not received for a generated message.

Event Destinations Security (Advanced)

Field	Description
MPModel	The message processing model to be used when generating SNMP messages using this entry.
SecurityModel	The security model to be used when generating SNMP messages using this entry.
SecurityName	Identifies the principal on whose behalf SNMP messages will be generated using this entry.
SecurityLevel	The level of security to be used when generating SNMP messages using this entry.

Event Filters General

Field	Description
FSPF - Nbr State Changes	Specifies whether or not the local switch should issue notification when the local switch learns of a change in the neighbor's state (state in the FSPF neighbor finite state machine) on an interface on a VSAN.
Domain Mgr - ReConfig Fabrics	Specifies whether or not the local switch should issue a notification on sending or receiving ReConfigureFabric (RCF) on a VSAN.
Zone Server - Request Rejects	Specifies if the zone server should issue a notification on rejects.

Field	Description
Zone Server - Merge Failures	Specifies if the zone server should issue a notification on merge failures.
Zone Server - Merge Successes	Specifies if the zone server should issue a notification on merge successes.
Zone Server - Default Zone Behavior Change	Specifies if the zone server should issue a notification if the propagation policy changes.
Zone Server - Unsupp Mode	Specifies if the zone server should issue a notification on unsupp mode changes
FabricConfigServer - Request Rejects	Specifies if the fabric configuration server should issue a notification on rejects.
RSCN - ILS Request Rejects	Specifies if the RSCN module should generate notifications when a SW_RSCN request is rejected.
RSCN - ILS RxRequest Rejects	Specifies if the RSCN module should generate notifications when a SW_RSCN request is rejected.
RSCN - ELS Request Rejects	Specifies if the RSCN module should generate notifications when a SCR or RSCN request is rejected.
FRU Changes	A false value will prevent field replaceable unit (FRU) notifications from being generated by this system.
SNMP - Community Auth Failure	Indicates whether the SNMP entity is permitted to generate authenticationFailure traps.
VRRP	Indicates whether the VRRP-enabled router will generate SNMP traps for events defined in this MIB.
FDMI	Specifies if the FDMI should generate notifications when a registration request is rejected.
License Manager	Indicates whether the system should generate notifications.
Port/Fabric Security	Specifies if the system should generate notifications when a port/fabric security issue arises.
FCC	Specifies whether the agent should generate notifications.
Name Server	If checked, the name server generates a notification when a request is rejected. If false, the notification is not generated.

Event Filters Interfaces

Field	Description
EnableLinkTrap	Indicates whether linkUp/linkDown traps should be generated for this interface.

Event Filters Control

Field	Description
Variable	Represents the notification to be controlled.
Descr	Description about the notification.
Enabled	Check to enable notification of the control. Shows the status of the control.

**Note**

You see the Descr column only on switches that run Cisco NX-OS Release 5.0 or later.



CHAPTER 51

Configuring System Message Logging

- [Configuring System Message Logging, on page 1023](#)

Configuring System Message Logging

This chapter describes how to configure system message logging on Cisco DCNM-SAN. It includes the following sections:

Information About System Message Logging

With the system message logging software, you can save messages in a log file or direct the messages to other devices. By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. This feature provides you with the following capabilities:

- Provides logging information for monitoring and troubleshooting
- Allows you to select the types of captured logging information
- Allows you to select the destination server to forward the captured logging information properly configured system message logging server.

You can monitor system messages by clicking the Events tab on DCNM-SAN or by choosing Logs > Events > Current on Device Manager. You can also monitor system messages remotely by accessing the switch through Telnet, SSH, or the console port, or by viewing the logs on a system message logging server.



Note

When the switch first initializes, the network is not connected until initialization completes. Therefore, messages are not redirected to a system message logging server for a few seconds.

Log messages are not saved across system reboots. However, a maximum of 100 log messages with a severity level of critical and below (levels 0, 1, and 2) are saved in NVRAM.

[Table 132: Internal Logging Facilities](#), on page 1024 describes some samples of the facilities supported by the system message logs.

Table 132: Internal Logging Facilities

Facility Keyword	Description	Standard or Cisco MDS Specific
acl	ACL manager	Cisco MDS 9000 Family specific
all	All facilities	Cisco MDS 9000 Family specific
auth	Authorization system	Standard
authpriv	Authorization (private) system	Standard
bootvar	Bootvar	Cisco MDS 9000 Family specific
callhome	Call Home	Cisco MDS 9000 Family specific
cron	Cron or at facility	Standard
daemon	System daemons	Standard
fcc	FCC	Cisco MDS 9000 Family specific
fcdomain	fcdomain	Cisco MDS 9000 Family specific
fcns	Name server	Cisco MDS 9000 Family specific
fcs	FCS	Cisco MDS 9000 Family specific
flogi	FLOGI	Cisco MDS 9000 Family specific
fspf	FSPF	Cisco MDS 9000 Family specific
ftp	File Transfer Protocol	Standard
ipconf	IP configuration	Cisco MDS 9000 Family specific
ipfc	IPFC	Cisco MDS 9000 Family specific
kernel	Kernel	Standard
local0 to local7	Locally defined messages	Standard
lpr	Line printer system	Standard
mail	Mail system	Standard
mcast	Multicast	Cisco MDS 9000 Family specific
module	Switching module	Cisco MDS 9000 Family specific
news	USENET news	Standard
ntp	NTP	Cisco MDS 9000 Family specific
platform	Platform manager	Cisco MDS 9000 Family specific

Facility Keyword	Description	Standard or Cisco MDS Specific
port	Port	Cisco MDS 9000 Family specific
port-channel	PortChannel	Cisco MDS 9000 Family specific
qos	QoS	Cisco MDS 9000 Family specific
rdl	RDL	Cisco MDS 9000 Family specific
rib	RIB	Cisco MDS 9000 Family specific
rscn	RSCN	Cisco MDS 9000 Family specific
securityd	Security	Cisco MDS 9000 Family specific
syslog	Internal system messages	Standard
sysmgr	System manager	Cisco MDS 9000 Family specific
tlport	TL port	Cisco MDS 9000 Family specific
user	User process	Standard
uucp	UNIX-to-UNIX Copy Program	Standard
vhbad	Virtual host base adapter daemon	Cisco MDS 9000 Family specific
vni	Virtual network interface	Cisco MDS 9000 Family specific
vrp_cfg	VRRP configuration	Cisco MDS 9000 Family specific
vrp_eng	VRRP engine	Cisco MDS 9000 Family specific
vsan	VSAN system messages	Cisco MDS 9000 Family specific
vshd	vshd	Cisco MDS 9000 Family specific
wwn	WWN manager	Cisco MDS 9000 Family specific
xbar	Xbar system messages	Cisco MDS 9000 Family specific
zone	Zone server	Cisco MDS 9000 Family specific

[Table 133: Error Message Severity Levels](#), on page 1025 describes the severity levels supported by the system message logs.

Table 133: Error Message Severity Levels

Level Keyword	Level	Description	System Message Definition
emergencies	0	System unusable	LOG_EMERG

Level Keyword	Level	Description	System Message Definition
alerts	1	Immediate action needed	LOG_ALERT
critical	2	Critical conditions	LOG_CRIT
errors	3	Error conditions	LOG_ERR
warnings	4	Warning conditions	LOG_WARNING
notifications	5	Normal but significant condition	LOG_NOTICE
informational	6	Informational messages only	LOG_INFO
debugging	7	Debugging messages	LOG_DEBUG

**Note**

Refer to the Cisco MDS 9000 Family and Nexus 7000 Series System Messages Reference for details on the error log message format.

This section includes the following topics:

Monitoring Syslog Server from DCNM-SAN

Cisco DCNM-SAN registers itself as a logging server and receives syslog messages and stores them in separate files for each switch.

With Cisco NX-OS Release 5.0(1a) and later, the DCNM-SAN stores the syslog messages from all the switches in a fabric to a database, and displays only the aggregated syslog information from the web client. This feature can be enabled or disabled. The syslog stored in the database is filtered by a configurable severity level.

Once the DCNM-SAN receives the syslog messages through the syslog receiver, the raw messages are parsed and the flag for persisting the message in the database is checked. The severity carried by this message is checked from the parsed fields, and the syslog messages are sent to the database.

The raw syslog messages are parsed into the following fields: switch time, facility, severity, event, and Vsan Id. The description is stored in the database and filtered by the severity level.

The following fields are added to server.properties:

- syslog.dblog.enable = false

This field is used to turn on the feature for storing the syslog messages into the database. By turning on this flag, the syslog messages are also written into the database.

- syslog.dblog.severity = warnings

This field is used to filter the syslog messages based on the severity. By configuring this property, syslog messages are filtered on the severity level.

System Message Logging

The system message logging software saves the messages in a log file or directs the messages to other devices. This feature has the following capabilities:

- Provides logging information for monitoring and troubleshooting.
- Allows the user to select the types of captured logging information.
- Allows the user to select the destination server to forward the captured logging information.

By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. You can specify which system messages should be saved based on the type of facility and the severity level. Messages are time-stamped to enhance real-time debugging and management.

You can access the logged system messages using the CLI or by saving them to a correctly configured system message logging server. The switch software saves system messages in a file that can save up to 1200 entries. You can monitor system messages remotely by accessing the switch through Telnet, SSH, the console port, or by viewing the logs on a system message logging server.

SFP Diagnostics

The error message related to SFP failures is written to the syslog. You can listen to the syslog for events related to SFP failures. The values, low or high alarm, and the warning are checked for the following parameters:

- TX Power
- RX Power
- Temperature
- Voltage
- Current

The SFP notification trap indicates the current status of the alarm and warning monitoring parameters for all the sensors based on the digital diagnostic monitoring information. This notification is generated whenever there is a change in the status of at least one of the monitoring parameters of the sensors on the transceiver in an interface.

The CISCO-INTERFACE-XCVR-MONITOR-MIB contains the SFP notification trap information. Refer to the Cisco MDS 9000 Family MIB Quick Reference for more information on this MIB.

Outgoing System Message Logging Server Facilities

All system messages have a logging facility and a level. The logging facility can be thought of as *where* and the level can be thought of as *what*.

The single system message logging daemon (syslogd) sends the information based on the configured **facility** option. If no facility is specified, local7 is the default outgoing facility.

The internal facilities are listed in [Table 132: Internal Logging Facilities](#), on page 1024 and the outgoing logging facilities are listed in [Table 134: Outgoing Logging Facilities](#), on page 1027.

Table 134: Outgoing Logging Facilities

Facility Keyword	Description	Standard or Cisco MDS Specific
auth	Authorization system	Standard
authpriv	Authorization (private) system	Standard
cron	Cron or at facility	Standard
daemon	System daemons	Standard

Facility Keyword	Description	Standard or Cisco MDS Specific
ftp	File Transfer Protocol	Standard
kernel	Kernel	Standard
local0 to local7	Locally defined messages	Standard (local7 is the default)
lpr	Line printer system	Standard
mail	Mail system	Standard
news	USENET news	Standard
syslog	Internal system messages	Standard
user	User process	Standard
uucp	UNIX-to-UNIX Copy Program	Standard

System Message Logging Servers

Device Manager allows you to view event logs on your local PC as well as those on the switch. For a permanent record of all events that occur on the switch, you should store these messages off the switch. To do this the Cisco MDS 9000 Family switch must be configured to send syslog messages to your local PC and a syslog server must be running on that PC to receive those messages. These messages can be categorized into four classes:

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



Note

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

System Message Logging Configuration Distribution

You can enable fabric distribution for all Cisco MDS 9000 Family switches in the fabric. When you perform system message logging configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you issue the first configuration command after you enabled distribution in a switch. The system message logging server uses the effective and pending database model to store or commit the commands based on your configuration. When you commit the configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. After making the configuration changes, you can

choose to discard the changes by aborting the changes instead of committing them. In either case, the lock is released. See [Chapter 13, “Using the CFS Infrastructure.”](#) for more information on the CFS application.

Fabric Lock Override

If you have performed a system message logging task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

Guidelines and Limitations

See the “CFS Merge Support” [section on page 13-5](#) for detailed concepts.

When merging two system message logging databases, follow these guidelines:

- Be aware that the merged database is a union of the existing and received database for each switch in the fabric.
- Verify that the merged database will only have a maximum of three system message logging servers.



Caution If the merged database contains more than three servers, the merge will fail.

Default Settings

[Table 135: Default System Message Log Settings](#), on [page 1029](#) lists the default settings for system message logging.

Table 135: Default System Message Log Settings

Parameters	Default
System message logging to the console	Enabled for messages at the critical severity level.
System message logging to Telnet sessions	Disabled.
Logging file size	4194304.
Log file name	Message (change to a name with up to 200 characters).
Logging server	Disabled.
Syslog server IP address	Not configured.
Number of servers	Three servers.
Server facility	Local 7.

Configuring System Message Logging

System logging messages are sent to the console based on the default (or configured) logging facility and severity values.

This sections includes the following topics:

Task Flow for Configuring System Message Logging

Follow these steps to configure system message logging:

Procedure

- Step 1** Enable or disable message logging.
 - Step 2** Configure console severity level.
 - Step 3** Configure monitor severity level.
 - Step 4** Configure module logging.
 - Step 5** Configure facility severity levels.
 - Step 6** Send log files.
 - Step 7** Configure system message logging servers.
 - Step 8** Configure system message logging distribution.
-

Enabling or Disabling Message Logging

You can disable logging to the console or enable logging to a specific Telnet or SSH session.

- When you disable or enable logging to a console session, that state is applied to all future console sessions. If you exit and log in again to a new session, the state is preserved.
- When you enable or disable logging to a Telnet or SSH session, that state is applied only to that session. If you exit and log in again to a new session, the state is not preserved.

To enable or disable the logging state for a Telnet or SSH session, follow these steps:

Procedure

- Step 1** Select a switch in the Fabric pane.
- Step 2** Expand **Events** and select **SysLog** in the Physical Attributes pane.
You see the SysLog information in the Information pane.
- Step 3** Click the **Switch Logging** tab.
You see the switch information.
- Step 4** Select a switch in the Information pane.
- Step 5** Check (enable) or uncheck (disable) the **Console Enable** check box.

- Step 6** Click the **Apply Changes** icon.
-

Configuring Console Severity Level

When logging is enabled for a console session (default), you can configure the severity levels of messages that appear on the console. The default severity for console logging is 2 (critical).

The current critical (default) logging level is maintained if the console baud speed is 9600 baud (default). All attempts to change the console logging level generates an error message. To increase the logging level (above critical), you must change the console baud speed to 38400 baud.

Configuring Monitor Severity Level

When logging is enabled for a monitor session (default), you can configure the severity levels of messages that appear on the monitor. The default severity for monitor logging is 5 (notifications).

To configure the severity level for a logging facility, follow these steps:

Procedure

- Step 1** Select a switch in the Fabric pane.
- Step 2** Expand **Events and** select **SysLog** in the Physical Attributes pane.
You see the SysLog information in the Information pane.
- Step 3** Click the **Switch Logging** tab.
You see the switch information.
- Step 4** Select a switch in the Information pane.
- Step 5** Select a severity level from the Console Severity drop-down list in the row for that switch.
- Step 6** Click the **Apply Changes** icon.
-

Configuring Module Logging

By default, logging is enabled at level 7 for all modules. You can enable or disable logging for each module at a specified level.

Configuring Facility Severity Levels

To configure the severity level for a logging facility, follow these steps:

Procedure

- Step 1** Expand **Events** and select **SysLog** in the Physical Attributes pane.
In Device Manager, choose **Logs > Syslog > Setup** and click the **Switch Logging** tab in the Syslog dialog box.

You see the switch information.

- Step 2** Check the check boxes where you want message logging to occur (**ConsoleEnable**, **TerminalEnable**, **LineCardEnable**).
- Step 3** Choose the message severity threshold from the Console Severity drop-down box for each switch in DCNM-SAN or click the appropriate message severity level radio button in Device Manager.
- Step 4** Click the Apply Changes icon in DCNM-SAN, or click Apply in Device Manager to save and apply your changes.

Sending Log Files

By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. Log messages are not saved across system reboots. The logging messages that are generated may be saved to a log file. You can configure the name of this file and restrict its size as required. The default log file name is messages.

The file name can have up to 80 characters and the file size ranges from 4096 bytes to 4194304 bytes.

The configured log file is saved in the /var/log/external directory. The location of the log file cannot be changed.

You can rename the log file using the **logging logfile** command.

The configured log file is saved in the /var/log/external directory. The location of the log file cannot be changed. You can use the **show logging logfile** and **clear logging logfile** commands to view and delete the contents of this file. You can use the **dir log:** command to view logging file statistics. You can use the **delete log:** command to remove the log file.

You can copy the logfile to a different location using the **copy log:** command using additional copy syntax.

To send log messages to a file, follow these steps:

Procedure

- Step 1** Select a switch in the Fabric pane.
- Step 2** Expand **Events** and select **SysLog** in the Physical Attributes pane.
You see the SysLog information in the Information pane.
- Step 3** Select a switch in the Information pane.
- Step 4** Click the **Switch Logging** tab.
- Step 5** Enter the name of the log file in the LogFile Name column in the row for that switch.
- Step 6** Click the **Apply Changes** icon.

Configuring System Message Logging Servers

You can configure a maximum of three system message logging servers. One of these syslog servers should be DCNM-SAN if you want to view system messages from the Event tab in DCNM-SAN.

To send log messages to a UNIX system message logging server, you must configure the system message logging daemon on a UNIX server. Log in as root, and follow these steps:

Procedure

Step 1 Add the following line to the `/etc/syslog.conf` file.

```
local1.debug/var/log/
myfile
.log
```

Note Be sure to add five tab characters between **local1.debug** and **/var/log/myfile.log**. Refer to entries in the `/etc/syslog.conf` file for further examples.

The switch sends messages according to the specified facility types and severity levels. The **local1** keyword specifies the UNIX logging facility used. The messages from the switch are generated by user processes. The **debug** keyword specifies the severity level of the condition being logged. You can set UNIX systems to receive all messages from the switch.

Step 2 Create the log file by entering these commands at the UNIX shell prompt:

```
$ touch /var/log/
myfile
.log
$ chmod 666 /var/log/
myfile
.log
```

Step 3 Make sure the system message logging daemon reads the new changes by entering this command:

```
$ kill -HUP ~cat /etc/syslog.pid~
```

What to do next



Note Most tabs in the Information pane for features using CFS are dimmed until you click the CFS tab. The CFS tab shows which switches have CFS enabled and shows the master switch for this feature. Once you click the CFS tab, the other tabs in the Information pane that use CFS are activated.

Configuring System Message Logging Servers

To configure system message logging servers, follow these steps:

Procedure

Step 1 Expand **Events** and select **SysLog** in the Physical Attributes pane.

Step 2 Click the Servers tab in the Information pane.

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

Step 3 Click the Create Row icon in DCNM-SAN, or click Create in Device Manager to add a new syslog server.

- Step 4** Enter the name or IP address in dotted decimal notation (for example, 192.168.2.12) of the syslog server in the Name or IP Address field.
 - Step 5** Set the message severity threshold by clicking the MsgSeverity radio button and set the facility by clicking the **Facility** radio button.
 - Step 6** Click the Apply Changes icon in DCNM-SAN, or click Create in Device Manager to save and apply your changes.
-

Fabric Lock Override

To use administrative privileges and release a locked system message logging session, use the **clear logging session** command.

```
switch# clear logging session
```

Verifying Syslog Servers from DCNM-SAN Web Server

To verify the syslog servers remotely using DCNM-SAN Web Server, follow these steps:

Procedure

- Step 1** Point your browser at the DCNM-SAN Web Server.
 - Step 2** Choose **Events > Syslog** to view the syslog server information for each switch. The columns in the table are sortable.
-

Monitoring Logs

This section covers the following topics:

Viewing Logs from DCNM-SAN Web Server

To view system messages remotely using DCNM-SAN Web Server, follow these steps:

Procedure

- Step 1** Point your browser at the DCNM-SAN Web Server.
 - Step 2** Click the **Events** tab followed by the Details to view the system messages. The columns in the events table are sortable. In addition, you can use the Filter button to limit the scope of messages within the table.
-

Viewing Logs from Device Manager

You can view system messages from Device Manager if Device Manager is running from the same workstation as DCNM-SAN. 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.



CHAPTER 52

Scheduling Maintenance Jobs

- [Scheduling Maintenance Jobs, on page 1037](#)

Scheduling Maintenance Jobs

The Cisco MDS command scheduler feature helps you schedule configuration and maintenance jobs in any switch in the Cisco MDS 9000 Family. You can use this feature to schedule jobs on a one-time basis or periodically.

This chapter includes the following sections:

Information About the Command Scheduler

The Cisco NX-OS command scheduler provides a facility to schedule a job (set of CLI commands) or multiple jobs at a specified time in the future. The job(s) can be executed once at a specified time in the future or at periodic intervals.

You can use this feature to schedule zone set changes, make QoS policy changes, back up data, save the configuration and do other similar jobs.

Scheduler Terminology

The following terms are used in this chapter:

- **Job**—A job is a set of NX-OS CLI commands (EXEC and config mode) that are executed as defined in the schedule.
- **Schedule**—A schedule determines the time when the assigned jobs must be executed. Multiple jobs can be assigned to a schedule. A schedule executes in one of two modes: one-time or periodic.
- **Periodic mode**—A job is executed at the user-specified periodic intervals, until it is deleted by the administrator. The following types of periodic intervals are supported:
 - **Daily**—The job is executed once a day.
 - **Weekly**—The job is executed once a week.
 - **Monthly**—The job is executed once a month.
 - **Delta**—The job is executed beginning at the specified start time and thereafter at user-specified intervals (days:hours:minutes).
- **One-time mode**—The job is executed once at a user-specified time.

Licensing Requirements for Command Scheduler

To use the command scheduler, you do not need to obtain any license.

Guidelines and Limitations

Before scheduling jobs on a Cisco MDS 9000 switch, note the following guidelines:

- Prior to Cisco MDS SAN-OS Release 3.0(3), only users local to the switch could perform scheduler configuration. As of Cisco MDS SAN-OS Release 3.0(3), remote users can perform job scheduling using AAA authentication.
- Be aware that the scheduled job can fail if it encounters one of the following situations when executing the job:
 - If the license has expired for a feature at the time when a job containing commands pertaining to that feature is scheduled.
 - If a feature is disabled at the time when a job containing commands pertaining to that feature is scheduled.
 - If you have removed a module from a slot and the job has commands pertaining to the interfaces for that module or slot.
- Verify that you have configured the time. The scheduler does not have any default time configured. If you create a schedule and assign job(s) and do not configure the time, that schedule is not launched.
- While defining a job, verify that no interactive or disruptive commands (for example, **copy bootflash:fileftp:URI**, **write erase**, and other similar commands) are specified as part of a job because the job is executed noninteractively at the scheduled time.

Default Settings

[Table 136: Default Command Scheduler Parameters](#), on page 1038 lists the default settings for command scheduling parameters.

Table 136: Default Command Scheduler Parameters

Parameters	Default
Command scheduler	Disabled.
Log file size	16 KB.



CHAPTER 53

Configuring RMON

- [Configuring RMON, on page 1039](#)

Configuring RMON

RMON is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. You can use the RMON alarms and events to monitor Cisco MDS 9000 Family switches running the Cisco SAN-OS Release 2.0(1b) or later or Cisco NX-OS Release 4.1(3) or later software.

This chapter includes the following sections:

Information About RMON

RMON is disabled by default, and no events or alarms are configured in the switch.

All switches in the Cisco MDS 9000 Family support the following RMON functions (defined in RFC 2819):

- **Alarm**—Each alarm monitors a specific management information base (MIB) object for a specified interval. When the MIB object value exceeds a specified value (rising threshold), the alarm condition is set and only one event is triggered regardless of how long the condition exists. When the MIB object value falls below a certain value (falling threshold), the alarm condition is cleared. This allows the alarm to trigger again when the rising threshold is crossed again.
- **Event**—Determines the action to take when an event is triggered by an alarm. The action can be to generate a log entry, an SNMP trap, or both.

For agent and management information, see the *Cisco MDS 9000 Family MIB Quick Reference*.

For information on an SNMP-compatible network management station, see the [Configuring RMON, on page 1039](#).

RMON Configuration Information

RMON is disabled by default and no events or alarms are configured in the switch. You can configure your RMON alarms and events by using the CLI or an SNMP-compatible network management station.



Tip

We recommend an additional, generic RMON console application on the network management station (NMS) to take advantage of RMON's network management capabilities.

RMON Configuration Using Threshold Manager

RMON is disabled by default and no events or alarms are configured in the switch. You can configure your RMON alarms and events by using the CLI or by using Threshold Manager in Device Manager.

The Threshold Monitor allows you to trigger an SNMP event or log a message when the selected statistic goes over a configured threshold value. RMON calls this a rising alarm threshold. The configurable settings are as follows:

- **Variable**—The statistic you want to set the threshold value on.
- **Value**—The value of the variable that you want the alarm to trigger at. This value is the difference (delta) between two consecutive polls of the variable by Device Manager.
- **Sample**—The sample period (in seconds) between two consecutive polls of the variable. Select your sample period such that the variable does not cross the threshold value you set under normal operating conditions.
- **Warning**—The warning level used by Device Manager to indicate the severity of the triggered alarm. This is a DCNM-SAN and Device Manager enhancement to RMON.



Note To configure any type of RMON alarm (absolute or delta, rising or falling threshold) click More on the Threshold Manager dialog box. You should be familiar with how RMON defines these concepts before configuring these advanced alarm types. Refer to the RMON-MIB (RFC 2819) for information on how to configure RMON alarms.



Note You must also configure SNMP on the switch to access RMON MIB objects.

RMON Alarm Configuration Information

Threshold Manager provides a list of common MIB objects to set an RMON threshold and alarm on. The alarm feature monitors a specific MIB object for a specified interval, triggers an alarm at a specified value (rising threshold), and resets the alarm at another value (falling threshold).

You can also set an alarm on any MIB object. The specified MIB must be an existing SNMP MIB object in standard dot notation (1.3.6.1.2.1.14.16777216 for ifInOctets.16777216).

Use one of the following options to specify the interval to monitor the MIB variable (ranges from 1 to 4294967295 seconds):

- Use the **delta** option to test the change between samples of a MIB variable.
- Use the **absolute** option to test each MIB variable directly.
- Use the **delta** option to test any MIB objects that are counters.

The range for the **rising threshold** and **falling threshold** values is -2147483647 to 2147483647.



Caution The **falling threshold** must be less than the **rising threshold**.

You can optionally specify the following parameters:

- The event-number to trigger if the rising or falling threshold exceeds the specified limit.

- The owner of the alarm.

Default Settings

Table 137: Default RMON Settings , on page 1041 lists the default settings for all RMON features in any switch.

Table 137: Default RMON Settings

Parameters	Default
RMON alarms	Disabled
RMON events	Disabled

Configuring RMON

RMON is disabled by default, and no events or alarms are configured in the switch.

This section includes the following topics:

Configuring the RMON Traps in SNMP

You must enable the RMON traps in the SNMP configuration for the RMON configuration to function correctly.



Note You must also configure SNMP on the switch to access RMON MIB objects.

Enabling RMON Alarms by Port

To configure an RMON alarm for one or more ports, follow these steps:

Procedure

- Step 1** Choose **Admin > Events > Threshold Manager** and click the FC Interfaces tab.
You see the Threshold Manager dialog box.
- Step 2** Choose the Select radio button to select individual ports for this threshold alarm.
 - a) Click the ... button to the right of the Selected field to display all ports.
 - b) Select the ports you want to monitor.
 - c) Click OK to accept the selection.

Alternatively, click the appropriate radio button to choose ports by type: All ports, xE ports, or Fx ports.
- Step 3** Check the check box for each variable to be monitored.
- Step 4** Enter the threshold value in the Value column.
- Step 5** Enter the sampling period in seconds. This is the time between each snapshot of the variable.

- Step 6** Choose one of the following severity levels to assign to the alarm: **Fatal, Warning, Critical, Error, Information.**
 - Step 7** Click Create.
 - Step 8** Confirm the operation to define an alarm and a log event when the system prompts you to define a severity event. If you do not confirm the operation, the system only defines a log event.
 - Step 9** Click More and then click the Alarms tab from the Threshold Manager dialog box to verify the alarm you created.
 - Step 10** Close both dialog box pop-up windows.
-

Enabling 32-Bit and 64-Bit Alarms

To configure an RMON alarm for one or more ports, follow these steps:

Procedure

- Step 1** Choose **Admin > Events > Threshold Manager** and click the **FC Interfaces > Create** tab.
You see the create 32-bit and 64-bit alarm dialog box.
 - Step 2** Click the **Select** radio button to select individual ports for this threshold alarm.
 - a) Click the **...** button to the right of the **Selected** field to display all ports.
 - b) Select the ports you want to monitor.
 - c) Click **OK** to accept the selection.

Alternatively, click the appropriate radio button to choose ports by type: **All ports, xE ports, or Fx ports.**
 - Step 3** Check the check box for each variable to be monitored.
 - Step 4** Enter the threshold value in the **Value** column.
 - Step 5** Enter the sampling period in seconds. This is the time between each snapshot of the variable.
 - Step 6** Choose one of the following severity levels to assign to the alarm: **Fatal, Warning, Critical, Error, Information.**
 - Step 7** Click Create.
 - Step 8** Confirm the operation to define an alarm and a log event when the system prompts you to define a severity event. If you do not confirm the operation, the system only defines a log event.
 - Step 9** Click More and then click the Alarms tab from the Threshold Manager dialog box to verify the alarm you created. The 32-bit and 64-bit alarm **Interval** columns show second as the unit.
 - Step 10** Close both dialog box pop-up windows.
-

Creating RMON Alarms

To create 64-bit RMON alarms, follow these steps:

Procedure

-
- Step 1** Expand Events and choose RMON from the **Physical Attributes pane**.
You see the 64-bit alarm dialog box.
- Step 2** Click the 64-bit alarms tab.
- Step 3** Click the Create Row tab. You see the Create Row window.
- Step 4** From the drop-down menu in the Variable field, choose from the list of MIB variables provided by the Threshold Manager.
- Note** You need to supply the interface details along with variables selected from the drop-down list to complete the Variable field, for example, ifHCInOctets.
- Step 5** Click the 32-bit alarms tab.
- Step 6** Click the Create Row tab.
- Step 7** From the drop-down menu in the Variable field, choose from the list of MIB variables provided by the Threshold Manager.
- Step 8** Click the radio button to choose the RMON alarm to be created (32-bit or 64-bit HC Alarm).
-

Enabling 32-Bit RMON Alarms for VSANs

To enable an RMON alarm for one or more VSANs, follow these steps:

Procedure

-
- Step 1** Choose **Admin > Events > Threshold Manager** and click the Services tab.
You see the Threshold Manager dialog box.
- Step 2** Click the **Services** tab.
You see the Threshold Manager dialog box with the Services tab for 32-bit alarm selected.
- Step 3** Click the 32-bit radio button.
- Step 4** Enter one or more VSANs (multiple VSANs separated by commas) to monitor in the VSAN ID(s) field. Use the down arrow to see a list of available VSANs to choose from.
- Step 5** Check the check box in the Select column for each variable to monitor.
- Step 6** Enter the threshold value in the Value column.
- Step 7** Enter the sampling period in seconds.
- Step 8** Choose a severity level to assign to the alarm: **Fatal, Critical, Error, Warning, Information**.
- Step 9** Click Create.
- Step 10** Confirm the operation to define an alarm and a log event when the system prompts you to define a severity event.
If you do not confirm the operation, the system only defines a log event.

- Step 11** Click More, and then click the Alarms tab in the Threshold Manager dialog box to verify the alarm you created.
-

Enabling 32-Bit and 64-Bit RMON Alarms for Physical Components

To configure an RMON alarm for a physical component for a 64-bit alarm, follow these steps:

Procedure

- Step 1** Choose **Admin** > Events > Threshold Manager and click the Physical tab.
You see the Threshold Manager dialog box with the Physical tab for the 64-bit alarm selected.
- Step 2** Check the check box in the Select column for each variable to monitor.
- Step 3** Enter the threshold value in the Value column.
- Step 4** Enter the sampling period in seconds.
- Step 5** Choose one of the following severity levels to assign to the alarm: Fatal(1), Warning(2), Critical(3), Error(4), Information(5).
- Step 6** Click Create.
- Step 7** Confirm the operation to define an alarm and a log event when the system prompts you to define a severity event.
If you do not confirm the operation, the system only defines a log event.
- Step 8** Click More, and then click the 64-bit Alarms tab in the Threshold Manager dialog box to verify the alarm you created.
- Note** The MaxAlarm option is noneditable because of backend support. The max RMON alarms cannot be set using the CLI.
-

Creating a New RMON from Device Manager Threshold Manager

RMON does not check the RMON alarm configuration before configuring the switch.

To configure an RMON alarm from Device Manager Threshold Manager, follow these steps:

Procedure

Expand Events, choose RMON, and click the Control tab.

You see the create RMON alarm Threshold Manager dialog box.

A user error is prompted if adding the new alarm exceeds the maximum alarm.

- Note** This feature is applicable when managing switches Release 4.1(1b) and later. Device Manager can only treat the existing alarm number as 0 for the checking.
-

Managing RMON Events

To define customized RMON events, follow these steps:

Procedure

- Step 1** Choose **Admin > Events > Threshold Manager** and click More in the Threshold Manager dialog box.
 - Step 2** Click the Events tab in the RMON Thresholds dialog box.
You see the RMON Thresholds Events tab.
 - Step 3** Click **Create** to create an event entry.
You see the Create RMON Thresholds Events dialog box.
 - Step 4** Configure the RMON threshold event attributes by choosing the type of event (**log**, **snmptrap**, or **logandtrap**).
 - Step 5** Increment the index. If you try to create an event with the existing index, you see a duplicate entry error message.
 - Step 6** (Optional) Provide a description and a community.
 - Step 7** Click Create, then close this dialog box.
 - Step 8** Verify that your event is listed in the remaining RMON Thresholds dialog box.
 - Step 9** Click Close to close the RMON Thresholds dialog box.
-

Managing RMON Alarms

To view the alarms that have already been enabled, follow these steps:

Procedure

- Step 1** Choose **Admin > Events > Threshold Manager** and click More in the Threshold Manager dialog box.
 - Step 2** Click the Alarms tab.
You see the RMON Thresholds dialog box.
 - Step 3** Delete any alarm by selecting it, and then click **Delete**.
-

Viewing the RMON Log

To view the RMON log, follow these steps:

Procedure

- Step 1** Choose **Admin > Events > Threshold Manager** and click More on the Threshold Manager dialog box.
- Step 2** Click the Log tab in the RMON Thresholds dialog box.

You see the RMON Thresholds Log tab. This is the log of RMON events that have been triggered by the Threshold Manager.

Field Descriptions for RMON

This section describes the field descriptions for RMON.

RMON Thresholds Controls

Field	Description
AlarmEnable	If true, the RMON alarm feature is enabled. If the RMON feature is disabled, all the RMON alarm related polling are stopped. Note that this is only intended for temporary disabling of RMON alarm feature to ensure that the CPU usage by RMON alarms is not detrimental. For permanent disabling on this feature, it suggested that all the entries in the alarmTable are removed.
MaxAlarms	The maximum number of entries allowed in the alarmTable.

RMON Thresholds 64bit Alarms

Field	Description
Interval	The interval in seconds over which the data is sampled and compared with the rising and falling thresholds. When setting this variable, care should be taken in the case of deltaValue sampling - the interval should be set short enough that the sampled variable is very unlikely to increase or decrease by more than $2^{31} - 1$ during a single sampling interval.
Variable	The variable to be sampled. Only variables that resolve to an ASN.1 primitive type of INTEGER (INTEGER, Integer32, Counter32, Counter64, Gauge, or TimeTicks) may be sampled.
SampleType	The method of sampling the selected variable and calculating the value to be compared against the thresholds. If the value is absoluteValue, the value of the selected variable will be compared directly with the thresholds at the end of the sampling interval. If the value is deltaValue, the value of the selected variable at the last sample will be subtracted from the current value, and the difference compared with the thresholds.
Value	The value of the statistic during the last sampling period. For example, if the sample type is deltaValue, this value will be the difference between the samples at the beginning and end of the period. If the sample type is absoluteValue, this value will be the sampled value at the end of the period. This is the value that is compared with the rising and falling thresholds. The value during the current sampling period is not made available until the period is completed and will remain available until the next period completes.
StartupAlarm	The alarm that may be sent when this entry is first set to valid.
Rising Threshold	A threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single event will be generated.

Field	Description
Rising EventId	The ID of the eventEntry that is used when a rising threshold is crossed.
Falling Threshold	A threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single event will be generated.
Falling EventId	The ID of the eventEntry that is used when a falling threshold is crossed. The eventEntry identified by a particular value of this index is the same as identified by the same value of eventIndex. If there is no corresponding entry in the eventTable, then no association exists. In particular, if this value is N/A, no associated event will be generated, as N/A is not a valid event index.
FailedAttempts	The number of times the alarm variable was polled (in the active state) and no response was received.
Owner	The ID of the user who configured this entry.

RMON Thresholds 32bit Alarms

Field	Description
Interval	The interval in seconds over which the data is sampled and compared with the rising and falling thresholds. When setting this variable, care should be taken in the case of deltaValue sampling - the interval should be set short enough that the sampled variable is very unlikely to increase or decrease by more than $2^{31} - 1$ during a single sampling interval.
Variable	The variable to be sampled. Only variables that resolve to an ASN.1 primitive type of INTEGER (INTEGER, Integer32, Counter32, Counter64, Gauge, or TimeTicks) may be sampled.
SampleType	The method of sampling the selected variable and calculating the value to be compared against the thresholds.
Value	The value of the statistic during the last sampling period.
StartupAlarm	The alarm that may be sent when this entry is first set to valid.
Rising Threshold	A threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single event will be generated.
Rising EventId	The ID of the eventEntry that is used when a rising threshold is crossed.
Falling Threshold	A threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single event will be generated.
Falling EventId	The ID of the eventEntry that is used when a falling threshold is crossed.
FailedAttempts	The number of times the alarm variable was polled (in the active state) and no response was received.

Field	Description
Owner	The ID of the user who configured this entry.

RMON Thresholds Events

Field	Description
Description	A comment describing this event entry.
Type	The type of notification that the probe will make about this event. In the case of log, an entry is made in the log table for each event. In the case of SNMP-trap, an SNMP trap is sent to one or more management stations.
LastTimeSent	When this event entry last generated an event. If this entry has not generated any events, this value will be N/A.
Owner	The entity that configured this entry and is therefore using the resources assigned to it.

RMON Thresholds Log

Field	Description
Time	When this log entry was created.
Description	A description of the event that activated this log entry.



CHAPTER 54

Configuring Fabric Configuration Server

- [Configuring Fabric Configuration Server, on page 1049](#)

Configuring Fabric Configuration Server

This chapter describes the Fabric Configuration Server (FCS) feature provided in the Cisco MDS 9000 Family of directors and switches. It includes the following sections:

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, Fx, and TL ports) and their attached Nx ports.
- Platform object—A set of nodes may be defined as a platform object to make it a single manageable entity. These nodes are end-devices (host systems, storage subsystems) attached to the fabric. Platform objects reside at the edge switches of the fabric.

Each object has its own set of attributes and values. A null value may also be defined for some attributes.

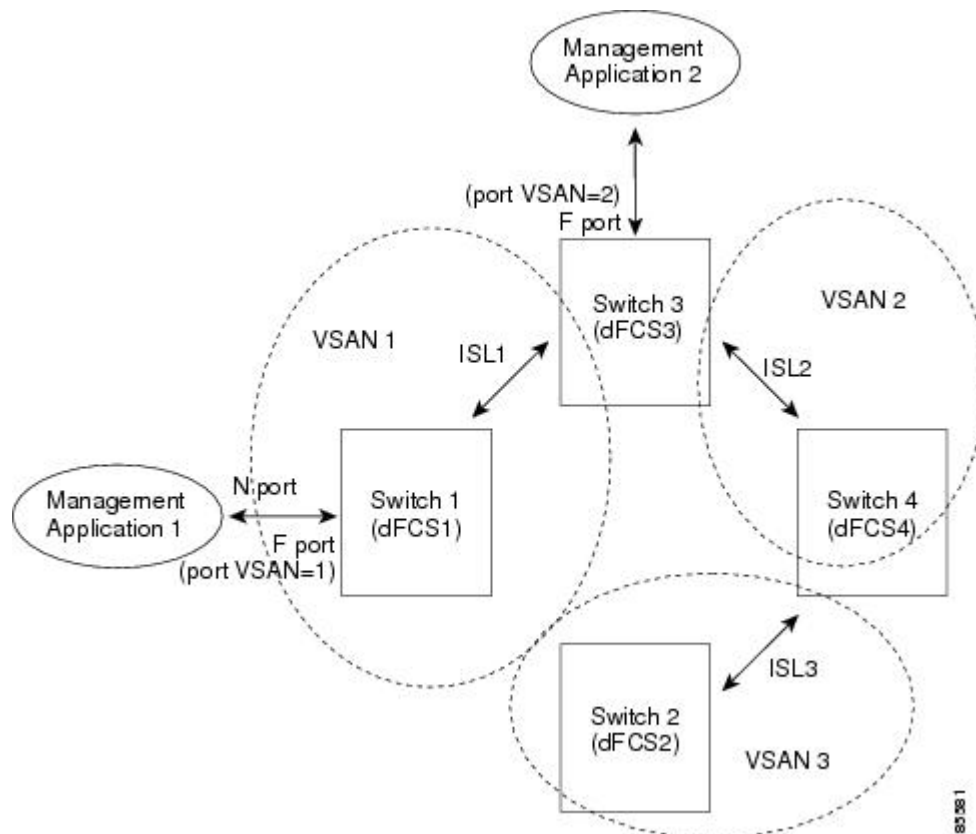
In the Cisco MDS 9000 Family switch environment, multiple VSANs constitute a fabric, where one instance of the FCS is present per VSAN.

As of Cisco NX-OS Release 4.1(1), FCS supports the discovery of virtual devices. The **fcs virtual-device-add** command, issued in FCS configuration submode, allows you to discover virtual devices in a particular VSAN or in all VSANs. The devices that are zoned for IVR must be discovered with this command and have request domain_ID (RDI) enabled, before activating the IVR zone set.

If you have attached a management application to a switch, all the frames directed towards the FCS in the switch are part of the port VSAN in the switch port (Fx port). Your view of the management application is limited only to this VSAN. However, information about other VSANs that this switch is part of can be obtained either through the SNMP or CLI.

In [Figure 166: FCSs in a VSAN Environment, on page 1050](#), Management Application 1 (M1) is connected through an F port with port VSAN ID 1, and Management Application 2 (M2) is connected through an F port with port VSAN ID 2. M1 can query the FCS information of switches S1 and S3, and M2 can query switches S3 and S4. Switch S2 information is not known to both of them. FCS operations can be done only on those switches that are visible in the VSAN. Note that M2 can send FCS requests only for VSAN 2 even though S3 is also a part of VSAN 1.

Figure 166: FCSs in a VSAN Environment



Significance of FCS

This section lists the significance of FCSs.

- FCSs support network management including the following:
 - N port management application can query and obtain information about fabric elements.
 - SNMP manager can use the FCS management information base (MIB) to start discovery and obtain information about the fabric topology.
- FCSs support TE and TL ports in addition to the standard F and E ports.
- FCS can maintain a group of modes with a logical name and management address when a platform registers with it. FCSs maintain a backup of all registrations in secondary storage and update it with every change. When a restart or switchover happens, FCSs retrieve the secondary storage information and rebuild its database.

- SNMP manager can query FCSs for all IEs, ports, and platforms in the fabric.

Default Settings

Table 138: Default FCS Settings , on page 1051 lists the default FCS settings.

Table 138: Default FCS Settings

Parameters	Default
Global checking of the platform name	Disabled.
Platform node type	Unknown.

Configuring FCS

The Fabric Configuration Server (FCS) provides discovery of topology attributes and maintains a repository of configuration information of fabric elements.

This section includes the following topic:

Specifying an FCS Name

You can specify if the unique name verification is for the entire fabric (globally) or only for locally (default) registered platforms.

Creating an FCS Platform

To create an FCS platform, follow these steps:

Procedure

-
- | | |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |
| Step 4 | Enter the VSAN ID, or select the ID from the drop-down list of available VSAN IDs. |
| Step 5 | Enter the Fabric Configuration Server name in the Name field. |
| Step 6 | Choose the type of server (Gateway , Host , Storage). |
| Step 7 | Enter the WWNs for the server. |
| Step 8 | Enter the management addresses for the server. |
| Step 9 | Click Create to create the server, or click Close to discard your changes and return to the Fabric Config Server dialog box. |
-

Displaying FCS Discovery

To display FCS discovery information, follow these steps:

Procedure

- Step 1** Choose **FC > Advanced > Fabric Config Server**.
You see the 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 Interconnect Element

To display FCS interconnect element information, follow these steps:

Procedure

- Step 1** Choose **FC > Advanced > Fabric Config Server**.
You see the Fabric Config Server dialog box.
- Step 2** Click the **Interconnect Elements** tab.
- Step 3** Click Close to close the dialog box.
-

Displaying FCS Fabric Ports

To display FCS discovery information, follow these steps:

Procedure

- 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.
- Step 3** Click **Refresh** to update the display.
-

Field Descriptions for FCS

This section displays the field descriptions for FCS.

Field	Description
FabricConfigServer - Request Rejects	Specifies if the Fabric Configuration Server should issue a notification on rejects.



CHAPTER 55

Monitoring Network Traffic Using SPAN

- [Monitoring Network Traffic Using SPAN, on page 1055](#)

Monitoring Network Traffic Using SPAN

This chapter describes the Switched Port Analyzer (SPAN) features provided in switches in the Cisco MDS 9000 Family.

This chapter includes the following sections:

Information About SPAN

The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic through a Fibre Channel interface. Traffic through any Fibre Channel interface can be replicated to a special port called the SPAN destination port (SD port). Any Fibre Channel port in a switch can be configured as an SD port. Once an interface is in SD port mode, it cannot be used for normal data traffic. You can attach a Fibre Channel Analyzer to the SD port to monitor SPAN traffic.

SD ports do not receive frames, they only transmit a copy of the SPAN source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports (see [Figure 167: SPAN Transmission , on page 1055](#)).

Figure 167: SPAN Transmission



This section covers the following topics:

SPAN Sources

SPAN sources refer to the interfaces from which traffic can be monitored. You can also specify VSAN as a SPAN source, in which case, all supported interfaces in the specified VSAN are included as SPAN sources. When a VSAN as a source is specified, then all physical ports and PortChannels in that VSAN are included as SPAN sources. You can choose the SPAN traffic in the ingress direction, the egress direction, or both directions for any source interface:

- Ingress source (Rx)—Traffic entering the switch fabric through this source interface is *spanned* or copied to the SD port (see [Figure 168: SPAN Traffic from the Ingress Direction , on page 1056](#)).

Figure 168: SPAN Traffic from the Ingress Direction

- Egress source (Tx)—Traffic exiting the switch fabric through this source interface is spanned or copied to the SD port (see [Figure 169: SPAN Traffic from Egress Direction, on page 1056](#)).

Figure 169: SPAN Traffic from Egress Direction

IPS Source Ports

SPAN capabilities are available on the IP Storage Services (IPS) module. The SPAN feature is only implemented on the FCIP and iSCSI virtual Fibre Channel port interfaces, not the physical Gigabit Ethernet ports. You can configure SPAN for ingress traffic, egress traffic, or traffic in both directions for all eight iSCSI and 24 FCIP interfaces that are available in the IPS module.



Note

You can configure SPAN for Ethernet traffic using Cisco switches or routers connected to the Cisco MDS 9000 Family IPS modules.

Allowed Source Interface Types

The SPAN feature is available for the following interface types:

- Physical ports such as F ports, FL ports, TE ports, E ports, and TL ports.
- Interface sup-fc0 (traffic to and from the supervisor):
 - The Fibre Channel traffic from the supervisor module to the switch fabric through the sup-fc0 interface is called ingress traffic. It is spanned when sup-fc0 is chosen as an ingress source port.
 - The Fibre Channel traffic from the switch fabric to the supervisor module through the sup-fc0 interface is called egress traffic. It is spanned when sup-fc0 is chosen as an egress source port.
- PortChannels
 - All ports in the PortChannel are included and spanned as sources.
 - You cannot specify individual ports in a PortChannel as SPAN sources. Previously configured SPAN-specific interface information is discarded.
- IPS module specific Fibre Channel interfaces:
 - iSCSI interfaces
 - FCIP interfaces

VSAN as a Source

SPAN sources refer to the interfaces from which traffic can be monitored. When a VSAN as a source is specified, then all physical ports and PortChannels in that VSAN are included as SPAN sources. A TE port is included only when the port VSAN of the TE port matches the source VSAN. A TE port is excluded even if the configured allowed VSAN list may have the source VSAN, but the port VSAN is different.

You cannot configure source interfaces (physical interfaces, PortChannels, or sup-fc interfaces) and source VSANs in the same SPAN session.

SPAN Sessions

Each SPAN session represents an association of one destination with a set of source(s) along with various other parameters that you specify to monitor the network traffic. One destination can be used by one or more SPAN sessions. You can configure up to 16 SPAN sessions in a switch. Each session can have several source ports and one destination port.

To activate any SPAN session, at least one source and the SD port must be up and functioning. Otherwise, traffic is not directed to the SD port.

**Tip**

A source can be shared by two sessions, however, each session must be in a different direction—one ingress and one egress.

You can temporarily deactivate (suspend) any SPAN session. The traffic monitoring is stopped during this time.

Specifying Filters

You can perform VSAN-based filtering to selectively monitor network traffic on specified VSANs. You can apply this VSAN filter to all sources in a session. Only VSANs present in the filter are spanned.

You can specify session VSAN filters that are applied to all sources in the specified session. These filters are bidirectional and apply to all sources configured in the session. Each SPAN session represents an association of one destination with a set of source(s) along with various other parameters that you specify to monitor the network traffic.

SD Port Characteristics

An SD port has the following characteristics:

- Ignores BB_credits.
- Allows data traffic only in the egress (Tx) direction.
- Does not require a device or an analyzer to be physically connected.
- Supports only 1 Gbps or 2 Gbps speeds. The auto speed option is not allowed.
- Multiple sessions can share the same destination ports.
- If the SD port is shut down, all shared sessions stop generating SPAN traffic.
- The outgoing frames can be encapsulated in Extended Inter-Switch Link (EISL) format.
- The SD port does not have a port VSAN.
- SD ports cannot be configured using Storage Services Modules (SSMs).
- The port mode cannot be changed if it is being used for a SPAN session.

**Note**

If you need to change an SD port mode to another port mode, first remove the SD port from all sessions and then change the port mode.

Monitoring Traffic Using Fibre Channel Analyzers

You can use SPAN to monitor traffic on an interface without any traffic disruption. This feature is especially useful in troubleshooting scenarios in which traffic disruption changes the problem environment and makes it difficult to reproduce the problem. You can monitor traffic in either of the following two ways:

- Without SPAN
- With SPAN

Monitoring Without SPAN

You can monitor traffic using interface fc1/1 in a Cisco MDS 9000 Family switch that is connected to another switch or host. You need to physically connect a Fibre Channel analyzer between the switch and the storage device to analyze the traffic through interface fc1/1 (see [Figure 170: Fibre Channel Analyzer Usage Without SPAN](#), on page 1058).

Figure 170: Fibre Channel Analyzer Usage Without SPAN



This type of connection has the following limitations:

- It requires you to physically insert the FC analyzer between the two network devices.
- It disrupts traffic when the Fibre Channel analyzer is physically connected.
- The analyzer captures data only on the Rx links in both port 1 and port 2. Port 1 captures traffic exiting interface fc1/1 and port 2 captures ingress traffic into interface fc1/1.

Monitoring with SPAN

Using SPAN you can capture the same traffic scenario without any traffic disruption. The Fibre Channel analyzer uses the ingress (Rx) link at port 1 to capture all the frames going out of the interface fc1/1. It uses the ingress link at port 2 to capture all the ingress traffic on interface fc1/1.

Using SPAN you can monitor ingress traffic on fc1/1 at SD port fc2/2 and egress traffic on SD port fc2/1. This traffic is seamlessly captured by the FC analyzer (see [Figure 171: Fibre Channel Analyzer Using SPAN](#), on page 1058).

Figure 171: Fibre Channel Analyzer Using SPAN



Single SD Port to Monitor Traffic

You do not need to use two SD ports to monitor bidirectional traffic on any interface. You can use one SD port and one FC analyzer port by monitoring traffic on the interface at the same SD port fc2/1.

[Figure 172: Fibre Channel Analyzer Using a Single SD Port](#), on page 1058 shows a SPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress directions. This setup is more advantageous and cost effective than the setup shown in the *Monitoring with SPAN* section. It uses one SD port and one port on the analyzer, instead of using a full, two-port analyzer.

Figure 172: Fibre Channel Analyzer Using a Single SD Port



To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

SD Port Configuration

The SD port in the destination switch enables the FC analyzer to receive the RSPAN traffic from the Fibre Channel tunnel. [Figure 173: RSPAN Tunnel Configuration, on page 1059](#) depicts an RSPAN tunnel configuration, now that tunnel destination is also configured.

Figure 173: RSPAN Tunnel Configuration



Note SD ports cannot be configured using Storage Services Modules (SSMs).

Mapping the FC Tunnel

The **tunnel-id-map** option specifies the egress interface of the tunnel at the destination switch (see [Figure 55-8](#)).



Creating VSAN Interfaces

[Figure 174: FC Tunnel Configuration, on page 1059](#) depicts a basic FC tunnel configuration.

Figure 174: FC Tunnel Configuration



Note This example assumes that VSAN 5 is already configured in the VSAN database.

Remote SPAN



Note Remote SPAN is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeSystem.

The Remote SPAN (RSPAN) feature enables you to remotely monitor traffic for one or more SPAN sources distributed in one or more source switches in a Fibre Channel fabric. The SPAN destination (SD) port is used for remote monitoring in a destination switch. A destination switch is usually different from the source switch(es) but is attached to the same Fibre Channel fabric. You can replicate and monitor traffic in any remote Cisco MDS 9000 Family switch or director, just as you would monitor traffic in a Cisco MDS source switch.

The RSPAN feature is nonintrusive and does not affect network traffic switching for those SPAN source ports. Traffic captured on the remote switch is tunneled across a Fibre Channel fabric which has trunking enabled

on all switches in the path from the source switch to the destination switch. The Fibre Channel tunnel is structured using trunked ISL (TE) ports. In addition to TE ports, the RSPAN feature uses two other interface types (see [Figure 175: RSPAN Transmission](#), on page 1060):

- SD ports—A passive port from which remote SPAN traffic can be obtained by the FC analyzer.
- ST ports—A SPAN tunnel (ST) port is an entry point port in the source switch for the RSPAN Fibre Channel tunnel. ST ports are special RSPAN ports and cannot be used for normal Fibre Channel traffic.

Figure 175: RSPAN Transmission



Advantages of Using RSPAN

The RSPAN features has the following advantages:

- Enables nondisruptive traffic monitoring at a remote location.
- Provides a cost effective solution by using one SD port to monitor remote traffic on multiple switches.
- Works with any Fibre Channel analyzer.
- Is compatible with the Cisco MDS 9000 Port Analyzer adapters.
- Does not affect traffic in the source switch, but shares the ISL bandwidth with other ports in the fabric.

FC and RSPAN Tunnels

An FC tunnel is a logical data path between a source switch and a destination switch. The FC tunnel originates from the source switch and terminates at the remotely located destination switch.

RSPAN uses a special Fibre Channel tunnel (FC tunnel) that originates at the ST port in the source switch and terminates at the SD port in the destination switch. You must bind the FC tunnel to an ST port in the source switch and map the same FC tunnel to an SD port in the destination switch. Once the mapping and binding is configured, the FC tunnel is referred to as an RSPAN tunnel (see [Figure 176: FC and RSPAN Tunnel](#), on page 1060).

Figure 176: FC and RSPAN Tunnel



ST Port Configuration

Once the FC tunnel is created, be sure to configure the ST port to bind it to the FC tunnel at the source switch. The FC tunnel becomes an RSPAN tunnel once the binding and mapping is complete.

[Figure 177: Binding the FC Tunnel](#), on page 1060 depicts a basic FC tunnel configuration.

Figure 177: Binding the FC Tunnel



ST Port Characteristics

ST ports have the following characteristics:

- ST ports perform the RSPAN encapsulation of the FC frame.
- ST ports do not use BB_credits.

- One ST port can only be bound to one FC tunnel.
- ST ports cannot be used for any purpose other than to carry RSPAN traffic.
- ST ports cannot be configured using Storage Services Modules (SSMs).

Creating Explicit Paths

You can specify an explicit path through the Cisco MDS Fibre Channel fabric (source-based routing), using the **explicit-path** option. For example, if you have multiple paths to a tunnel destination, you can use this option to specify the FC tunnel to always take one path to the destination switch. The software then uses this specified path even if other paths are available.

This option is especially useful if you prefer to direct the traffic through a certain path although other paths are available. In an RSPAN situation, you can specify the explicit path so the RSPAN traffic does not interfere with the existing user traffic. You can create any number of explicit paths in a switch (see [Figure 178: Explicit Path Configuration, on page 1061](#)).

Figure 178: Explicit Path Configuration



Guidelines and Limitations

SPAN Configuration Guidelines

The following guidelines and limitations apply for SPAN configurations:

- You can configure up to 16 SPAN sessions with multiple ingress (Rx) sources.
- You can configure a maximum of three SPAN sessions with one egress (Tx) port.
- In a 32-port switching module, you must configure the same session in all four ports in one port group (unit). If you wish, you can also configure only two or three ports in this unit.
- SPAN frames are dropped if the sum of the bandwidth of the sources exceeds the speed of the destination port.
- Frames dropped by a source port are not spanned.
- SPAN does not capture pause frames in a Fibre Channel over Ethernet (FCoE) network because pause frames sent from the virtual expansion (VE) port are generated and terminated by the outermost MAC layer. For more information on FCoE, see the Cisco NX-OS FCoE Configuration Guide for Cisco Nexus 7000 and Cisco MDS 9500.

Guidelines to Configure VSANs as a Source

The following guidelines apply when configuring VSANs as a source:

- Traffic on all interfaces included in a source VSAN is spanned only in the ingress direction.
- If a VSAN is specified as a source, you cannot perform interface-level SPAN configuration on the interfaces that are included in the VSAN. Previously configured SPAN-specific interface information is discarded.
- If an interface in a VSAN is configured as a source, you cannot configure that VSAN as a source. You must first remove the existing SPAN configurations on such interfaces before configuring VSAN as a source.
- Interfaces are only included as sources when the port VSAN matches the source VSAN. [Figure 179: VSAN as a Source , on page 1062](#) displays a configuration using VSAN 2 as a source:

- All ports in the switch are in VSAN 1 except fc1/1.
- Interface fc1/1 is the TE port with port VSAN 2. VSANs 1, 2, and 3 are configured in the allowed list.
- VSAN 1 and VSAN 2 are configured as SPAN sources.

Figure 179: VSAN as a Source

For this configuration, the following apply:

- VSAN 2 as a source includes only the TE port fc1/1 that has port VSAN 2.
- VSAN 1 as a source does not include the TE port fc1/1 because the port VSAN does not match VSAN 1.

Guidelines to Specifying Filters

The following guidelines apply to SPAN filters:

- PortChannel configurations are applied to all ports in the PortChannel.
- If no filters are specified, the traffic from all active VSANs for that interface is spanned by default.
- While you can specify arbitrary VSAN filters in a session, traffic can only be monitored on the port VSAN or on allowed-active VSANs in that interface.

RSPAN Configuration Guidelines

The following guidelines apply for a SPAN configuration:

- All switches in the end-to-end path of the RSPAN tunnel must belong to the Cisco MDS 9000 Family.
- All VSANs with RSPAN traffic must be enabled. If a VSAN containing RSPAN traffic is not enabled, it is dropped.
- The following configurations must be performed on *each* switch in the end-to-end path of the Fibre Channel tunnel in which RSPAN is to be implemented:
 - Trunking must be enabled (enabled by default) and the trunk enabled link must be the lowest cost link in the path.
 - VSAN interface must be configured.
 - The Fibre Channel tunnel feature must be enabled (disabled by default).
 - IP routing must be enabled (disabled by default).



Note

If the IP address is in the same subnet as the VSAN, the VSAN interface does not have to be configured for all VSANs on which the traffic is spanned.

- A single Fibre Channel switch port must be dedicated for the ST port functionality.
- Do not configure the port to be monitored as the ST port.
- The FC tunnel's IP address must reside in the same subnet as the VSAN interface.

Default SPAN and RSPAN Settings

[Table 139: Default SPAN Configuration Parameters](#), on page 1063 lists the default settings for SPAN parameters.

Table 139: Default SPAN Configuration Parameters

Parameters	Default
SPAN session	Active.
If filters are not specified	SPAN traffic includes traffic through a specific interface from all active VSANs.
Encapsulation	Disabled.
SD port	Output frame format is Fibre Channel.

[Table 140: Default RSPAN Configuration Parameters](#), on page 1063 lists the default settings for RSPAN parameters.

Table 140: Default RSPAN Configuration Parameters

Parameters	Default
FC tunnel	Disabled.
Explicit path	Not configured.
Minimum cost path	Used if explicit path is not configured.

Configuring SPAN

The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic through a Fibre Channel interface.

This section covers the following topics:

Configuring SD Ports for SPAN

To monitor network traffic using SD ports, follow these steps:

Procedure

-
- | | |
|---------------|---------------------------------------------------------------------|
| 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. |
-

Configuring SD Ports for SPAN using DM

To configure an SD port for SPAN monitoring using Device Manager, follow these steps:

Procedure

- Step 1** Right-click the port you want to configure and select **Configure**.
You see the general port configuration dialog.
- Step 2** Under Mode, choose **SD**.
- Step 3** Click **Apply** to accept the change.
- Step 4** Close the dialog box.
-

Configuring SPAN max-queued-packets

When a SPAN destination port is oversubscribed or has more source traffic than the speed of the destination port, the source ports of the SPAN session will reduce in their throughput. The impact is proportional to the amount of source traffic flowing in. Lowering the max-queued-packets value from the default value of 15 to 1 prevents the impact on the source ports. It is necessary to reconsider the default value for this setting as it may impact the source interface throughput.

The following are the requirements:

- The span max-queued-packets can be changed only if no SPAN sessions are currently active on the switch.
- If you are spanning the traffic going through an FCIP interface, SPAN copies may be dropped even if the SD interface has more bandwidth than the amount of traffic being replicated. To avoid SPAN drops, set the max-queued-packets to a higher value; for example, 100.

By default, SPAN frames are dropped if the sum of the bandwidth of the source interfaces exceed the bandwidth of the destination port. With a higher value, the SPAN traffic has a higher probability of reaching the SPAN destination port instead of being dropped at the expense of data traffic throughput.

Creating SPAN Sessions

To create SPAN sessions, follow these steps:

Procedure

- 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.
- Step 4** Choose the session ID (the ID range may vary depending on platform type and version) using the up or down arrows and click Create.
- Step 5** Repeat Step 4 for each session 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.

Configuring SPAN for Generation 2 Fabric Switches

Cisco Generation 2 fabric switches (such as MDS 9124) support SPAN sessions in both directions, Rx and Tx.



Note While using Generation 2 fabric switches, you cannot create an additional active SPAN session when you already have one.

the following are the restrictions:

- You can specify multiple SPAN source interfaces in Rx and Tx directions. However, the direction should be explicitly mentioned at the end of the command. The SPAN will reject any source interface configuration that fails to mention the direction.
- You cannot mix ingress and egress interfaces in the same SPAN session. The SPAN will reject any configuration that mixes Rx and Tx directions. However, you can specify multiple SPAN source interfaces in a single direction.

Editing SPAN Sources

To edit a SPAN source, follow these steps:

Procedure

- Step 1** Choose Interface > SPAN.
You see the SPAN dialog box.
- Step 2** Click the 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.
You see the 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** Click the direction (**receive** or **transmit**) 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

To delete a SPAN session, follow these steps:

Procedure

- Step 1** Choose Interface > SPAN.
You see the SPAN dialog box.
- Step 2** Click the Sessions tab.
- Step 3** Click the SPAN session you want to delete.
- Step 4** Click Delete.
The SPAN session is deleted.
- Step 5** Close the dialog box.
-

Encapsulating Frames

The frame encapsulation feature is disabled by default. If you enable the encapsulation feature, all outgoing frames are encapsulated.

Configuring Fibre Channel Analyzers Using SPAN

To configure Fibre Channel Analyzers using SPAN for the example in the *Fibre Channel Analyzer Using SPAN* section, follow these steps:

Procedure

- Step 1** Configure SPAN on interface fc1/1 in the ingress (Rx) direction to send traffic on SD port fc2/1 using session 1.
- Step 2** Configure SPAN on interface fc1/1 in the egress (Tx) direction to send traffic on SD port fc2/2 using session 2.
- Step 3** Physically connect fc2/1 to port 1 on the Fibre Channel analyzer.
- Step 4** Physically connect fc2/2 to port 2 on the Fibre Channel analyzer.
-

Configuring RSPAN

The RSPAN tunnel begins in the source switch and terminates in the destination switch. This section assumes Switch S to be the source and Switch D to be the destination.

The following are the prerequisites:

- In addition to the source and destination switches, the VSAN must also be configured in each Cisco MDS switch in the Fibre Channel fabric, if they exist.

To monitor network traffic using the RSPAN feature, follow these steps:

Procedure

-
- | | |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Create VSAN interfaces in destination switch (Switch D) and source switch (Switch S) to facilitate the Fibre Channel tunnel (FC tunnel) creation. |
| Step 2 | Enable the FC tunnel in each switch in the end-to-end path of the tunnel. |
| Step 3 | Initiate the FC tunnel (in Switch S) and map the tunnel to the VSAN interface's IP address (in Switch D) so all RSPAN traffic from the tunnel is directed to the SD port. |
| Step 4 | Configure SD ports for SPAN monitoring in the destination switch (Switch D). |
| Step 5 | Configure the ST port in the source switch (Switch S) and bind the ST port to the FC tunnel. |
| Step 6 | Create an RSPAN session in the source switch (in Switch S) to monitor network traffic. |
-

Configuring the Source Switch

This section identifies the tasks that must be performed in the source switch (Switch S):

Enabling FC Tunnels

The following are the restrictions:

- FC tunnels do not work over nontrunking ISLs.
- The interface cannot be operationally up until the FC tunnel mapping is configured in the destination switch.



Note Be sure to enable this feature in each switch in the end-to-end path in the fabric.

Configuring All Intermediate Switches

This section identifies the tasks that must be performed in all intermediate switches in the end-to-end path of the RSPAN tunnel:

Configuring VSAN Interfaces

An RSPAN tunnel configuration is terminated in the destination switch (Switch D).



Note This example assumes that VSAN 5 is already configured in the VSAN database.

Enabling IP Routing

The IP routing feature is disabled by default. Be sure to enable IP routing in each switch (including the source and destination switches) in the end-to-end path in the fabric. This procedure is required to set up the FC tunnel.

Configuring the Destination Switch

This section identifies the tasks that must be performed in the destination switch (Switch D):

Configuring the SD Port

SD ports cannot be configured using Storage Services Modules (SSMs).

Monitoring RSPAN Traffic

Once the session is configured, other SPAN sources for this session can also be configured as required. [Figure 180: Fibre Channel Analyzer Using a Single SD Port to Monitor RSPAN Traffic, on page 1068](#) shows an RSPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress directions.

Figure 180: Fibre Channel Analyzer Using a Single SD Port to Monitor RSPAN Traffic



To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

Configuration Examples for RSPAN

This section covers the following topic:



Note

RSPAN can be combined with the local SPAN feature so SD ports forward local SPAN traffic along with remote SPAN traffic. Various SPAN source and tunnel scenarios are described in this section.

Single Source with One RSPAN Tunnel

The source Switch S and the destination Switch D are interconnected through a Fibre Channel fabric. An RSPAN tunnel is configured as a destination interface for the SPAN session and the ST port forwards SPAN traffic through the RSPAN tunnel (see [Figure 181: RSPAN Scenario with One Source Switch, One Destination Switch, and One Tunnel, on page 1068](#)).

Figure 181: RSPAN Scenario with One Source Switch, One Destination Switch, and One Tunnel



Single Source with Multiple RSPAN Tunnels

[Figure 182: RSPAN Scenario with One Source Switch, One Destination Switch, and Multiple Tunnels, on page 1069](#) displays two separate RSPAN tunnels configured between Switches S and N. Each tunnel has an associated ST port in the source switch and a separate SD port in the destination switch. This configuration is useful for troubleshooting purposes.

Figure 182: RSPAN Scenario with One Source Switch, One Destination Switch, and Multiple Tunnels



Multiple Sources with Multiple RSPAN Tunnels

Figure 183: RSPAN Scenario with Two Source Switches, a Destination Switch, and Multiple Tunnels, on page 1069 displays two separate RSPAN tunnels configured between Switches S1 and S2. Both tunnels have an associated ST port in their respective source switch and terminate in the same SD port in the destination switch.

Figure 183: RSPAN Scenario with Two Source Switches, a Destination Switch, and Multiple Tunnels



This configuration is useful for remote monitoring purposes. For example, the administrator may be at the destination switch and can remotely monitor the two source switches.

Field Descriptions for SPAN

This section describes the field descriptions for SPAN.

SPAN Sessions

Field	Description
Dest Interface	The Span Destination port interface.
Filter VSAN List	The VSANs that are assigned to this session.
Status Admin	Suspend an active session or activate an inactive session.
Status Oper	The current state of the session.
Description	The description of the session status.
VSAN List	The VSANs that are assigned to this session.
Or Interface (Direction)	The destination port ID to be configured for the session.
Inactive Reason	Description of the reason why this session is not active.

Related Topics

[SPAN Sessions, on page 351](#)

[Creating SPAN Sessions, on page 358](#)

[Deleting SPAN Sessions, on page 360](#)

[Information About SPAN, on page 349](#)

[Editing SPAN Sources, on page 359](#)

Span Global

Field	Description
MaxQueuedSpanPackets	This field specifies the drop threshold packets for all span sessions. The MaxQueuedSpanPackets field is only available when no session is active.

SPAN Source Interfaces

Field	Description
Interface, Direction	The destination port ID configured for the session, and the direction of traffic.



CHAPTER 56

Monitoring System Processes and Logs

- [Monitoring System Processes and Logs, on page 1071](#)

Monitoring System Processes and Logs

This chapter provides details on monitoring the health of the switch and includes the following sections.

Information About System Processes and Logs

This section includes the following topics:

Saving Cores

You can save cores (from the active supervisor module, the standby supervisor module, or any switching module) to an external CompactFlash (slot 0) or to a TFTP server in one of two ways:

- On demand—Copies a single file based on the provided process ID.
- Periodically—Copies core files periodically as configured by the user.

A new scheme overwrites any previously issued scheme. For example, if you perform another core log copy task, the cores are periodically saved to the new location or file.

Saving the Last Core to Bootflash

This last core dump is automatically saved to bootflash in the /mnt/pss/ partition before the switchover or reboot occurs. Three minutes after the supervisor module reboots, the saved last core is restored from the flash partition (/mnt/pss) back to its original RAM location. This restoration is a background process and is not visible to the user.



Tip The timestamp on the restored last core file displays the time when the supervisor booted up not when the last core was actually dumped. To obtain the exact time of the last core dump, check the corresponding log file with the same PID.

To view the last core information, enter the **show cores** command in EXEC mode.

To view the time of the actual last core dump, enter the **show process log** command in EXEC mode.

First and Last Core

The first and last core feature uses the limited system resource and retains the most important core files. Generally, the first core and the most recently generated core have the information for debugging and, the first and last core feature tries to retain the first and the last core information.

If the core files are generated from an active supervisor module, the number of core files for the service is defined in the `service.conf` file. There is no upper limit on the total number of core files in the active supervisor module.

To display the core files saved in the system, use the `show cores` command.

Online System Health Management

The Online Health Management System (OHMS) (system health) is a hardware fault detection and recovery feature. It ensures the general health of switching, services, and supervisor modules in any switch in the Cisco MDS 9000 Family.

The OHMS monitors system hardware in the following ways:

- The OHMS component running on the active supervisor maintains control over all other OHMS components running on the other modules in the switch.
- The system health application running in the standby supervisor module only monitors the standby supervisor module, if that module is available in the HA standby mode.

The OHMS application launches a daemon process in all modules and runs multiple tests on each module to test individual module components. The tests run at preconfigured intervals, cover all major fault points, and isolate any failing component in the MDS switch. The OHMS running on the active supervisor maintains control over all other OHMS components running on all other modules in the switch.

On detecting a fault, the system health application attempts the following recovery actions:

- Performs additional testing to isolate the faulty component.
- Attempts to reconfigure the component by retrieving its configuration information from persistent storage.
- If unable to recover, sends Call Home notifications, system messages and exception logs; and shuts down and discontinues testing the failed module or component (such as an interface).
- Sends Call Home and system messages and exception logs as soon as it detects a failure.
- Shuts down the failing module or component (such as an interface).
- Isolates failed ports from further testing.
- Reports the failure to the appropriate software component.
- Switches to the standby supervisor module, if an error is detected on the active supervisor module and a standby supervisor module exists in the Cisco MDS switch. After the switchover, the new active supervisor module restarts the active supervisor tests.
- Reloads the switch if a standby supervisor module does not exist in the switch.
- Provides CLI support to view, test, and obtain test run statistics or change the system health test configuration on the switch.
- Performs tests to focus on the problem area.

Each module is configured to run the test relevant to that module. You can change the default parameters of the test in each module as required.

Loopback Test Configuration Frequency

Loopback tests are designed to identify hardware errors in the data path in the module(s) and the control path in the supervisors. One loopback frame is sent to each module at a preconfigured frequency—it passes through each configured interface and returns to the supervisor module.

The loopback tests can be run at frequencies ranging from 5 seconds (default) to 255 seconds. If you do not configure the loopback frequency value, the default frequency of 5 seconds is used for all modules in the switch. Loopback test frequencies can be altered for each module.

Loopback Test Configuration Frame Length

Loopback tests are designed to identify hardware errors in the data path in the module(s) and the control path in the supervisors. One loopback frame is sent to each module at a preconfigured size—it passes through each configured interface and returns to the supervisor module.

The loopback tests can be run with frame sizes ranging from 0 bytes to 128 bytes. If you do not configure the loopback frame length value, the switch generates random frame lengths for all modules in the switch (auto mode). Loopback test frame lengths can be altered for each module.

Hardware Failure Action

The failure-action command controls the Cisco NX-OS software from taking any action if a hardware failure is determined while running the tests.

By default, this feature is enabled in all switches in the Cisco MDS 9000 Family—action is taken if a failure is determined and the failed component is isolated from further testing.

Failure action is controlled at individual test levels (per module), at the module level (for all tests), or for the entire switch.

Performing Test Run Requirements

Enabling a test does not guarantee that the test will run.

Tests on a specific interface or module only run if you enable system health for all of the following items:

- The entire switch
- The required module
- The required interface

**Tip**

The test will not run if system health is disabled in any combination. If system health is disabled to run tests, the test status shows up as disabled.

**Tip**

If the specific module or interface is enabled to run tests, but is not running the tests due to system health being disabled, then tests show up as enabled (not running).

Tests for a Specified Module

The system health feature in the NX-OS software performs tests in the following areas:

- Active supervisor's in-band connectivity to the fabric.
- Standby supervisor's arbiter availability.
- Bootflash connectivity and accessibility on all modules.
- EOBC connectivity and accessibility on all modules.
- Data path integrity for each interface on all modules.
- Management port's connectivity.
- User-driven test for external connectivity verification, port is shut down during the test (Fibre Channel ports only).
- User-driven test for internal connectivity verification (Fibre Channel and iSCSI ports).

Clearing Previous Error Reports

You can clear the error history for Fibre Channel interfaces, iSCSI interfaces, an entire module, or one particular test for an entire module. By clearing the history, you are directing the software to retest all failed components that were previously excluded from tests.

If you previously enabled the failure-action option for a period of time (for example, one week) to prevent OHMS from taking any action when a failure is encountered and after that week you are now ready to start receiving these errors again, then you must clear the system health error status for each test.



Tip

The management port test cannot be run on a standby supervisor module.

Interpreting the Current Status

The status of each module or test depends on the current configured state of the OHMS test in that particular module (see [Table 141: OHMS Configured Status for Tests and Modules](#), on page 1074).

Table 141: OHMS Configured Status for Tests and Modules

Status	Description
Enabled	You have currently enabled the test in this module and the test is not running.
Disabled	You have currently disabled the test in this module.
Running	You have enabled the test and the test is currently running in this module.
Failing	This state is displayed if a failure is imminent for the test running in this module—possibility of test recovery exists in this state.
Failed	The test has failed in this module—and the state cannot be recovered.
Stopped	The test has been internally stopped in this module by the Cisco NX-OS software.
Internal failure	The test encountered an internal failure in this module. For example, the system health application is not able to open a socket as part of the test procedure.
Diags failed	The startup diagnostics has failed for this module or interface.
On demand	The system health external-loopback or the system health internal-loopback tests are currently running in this module. Only these two commands can be issued on demand.

Status	Description
Suspended	Only encountered in the MDS 9100 Series due to one oversubscribed port moving to a E or TE port mode. If one oversubscribed port moves to this mode, the other three oversubscribed ports in the group are suspended.

The status of each test in each module is visible when you display any of the **show system health** commands.

On-Board Failure Logging

The Generation 2 Fibre Channel switching modules provide the facility to log failure data to persistent storage, which can be retrieved and displayed for analysis. This on-board failure logging (OBFL) feature stores failure and environmental information in nonvolatile memory on the module. The information will help in post-mortem analysis of failed cards.

OBFL data is stored in the existing CompactFlash on the module. OBFL uses the persistent logging (PLOG) facility available in the module firmware to store data in the CompactFlash. It also provides the mechanism to retrieve the stored data.

The data stored by the OBFL facility includes the following:

- Time of initial power-on
- Slot number of the card in the chassis
- Initial temperature of the card
- Firmware, BIOS, FPGA, and ASIC versions
- Serial number of the card
- Stack trace for crashes
- CPU hog information
- Memory leak information
- Software error messages
- Hardware exception logs
- Environmental history
- OBFL specific history information
- ASIC interrupt and error statistics history
- ASIC register dumps

Default Settings

[Table 142: Default System Health and Log Settings](#), on page 1075 lists the default system health and log settings.

Table 142: Default System Health and Log Settings

Parameters	Default
Kernel core generation	One module
System health	Enabled
Loopback frequency	5 seconds
Failure action	Enabled

Core and Log Files

This section includes the following topics:

Clearing the Core Directory

Clearing the Core Directory

Use the **clear cores** command to clean out the core directory. The software clears all the core files and other cores present on the active supervisor module.

```
switch# clear cores
```

Before you begin

Ensure that SSH2 is enabled on this switch.

To clear the cores on a switch, follow these steps:

Procedure

Step 1 Click **Clear** to clear the cores.

The software keeps the last few cores, per service and per slot, and clears all the core files and other cores present on the active supervisor module.

Step 2 Click **Close** to close the dialog box.

Configuring System Health

The Online Health Management System (OHMS) (system health) is a hardware fault detection and recovery feature. It ensures the general health of switching, services, and supervisor modules in any switch in the Cisco MDS 9000 Family.

This section includes the following topics:

Task Flow for Configuring System Health

Follow these steps to configure system health:

Procedure

-
- Step 1** Enable System Health Initiation.
 - Step 2** Configure Loopback Test Configuration Frequency.
 - Step 3** Configure Loopback Test Configuration Frame Length.
 - Step 4** Configure Hardware Failure Action.
 - Step 5** Perform Test Run Requirements.

- Step 6** Clear Previous Error Reports.
- Step 7** Perform Internal Loopback Tests.
- Step 8** Perform External Loopback Tests.
- Step 9** Perform Serdes Loopbacks.

Performing Internal Loopback Tests

You can run manual loopback tests to identify hardware errors in the data path in the switching or services modules, and the control path in the supervisor modules. Internal loopback tests send and receive FC2 frames to and from the same ports and provide the round-trip time taken in microseconds. These tests are available for Fibre Channel, IPS, and iSCSI interfaces.



Note If the test fails to complete successfully, the software analyzes the failure and prints the following error:
External loopback test on interface fc 7/2 failed. Failure reason: Failed to loopback, analysis complete Failed device ID 3 on module 1

Choose **Interface > Diagnostics > Internal** to perform an internal loopback test from Device Manager.

Performing External Loopback Tests

You can run manual loopback tests to identify hardware errors in the data path in the switching or services modules, and the control path in the supervisor modules. External loopback tests send and receive FC2 frames to and from the same port or between two ports.

You need to connect a cable (or a plug) to loop the Rx port to the Tx port before running the test. If you are testing to and from the same port, you need a special loop cable. If you are testing to and from different ports, you can use a regular cable. This test is only available for Fibre Channel interfaces.



Note If the test fails to complete successfully, the software analyzes the failure and prints the following error:
External loopback test on interface fc 7/2 failed. Failure reason: Failed to loopback, analysis complete Failed device ID 3 on module 1

Choose **Interface > Diagnostics > External** to perform an external loopback test from Device Manager.

Performing Serdes Loopbacks

Serializer/Deserializer (serdes) loopback tests the hardware for a port. These tests are available for Fibre Channel interfaces.



Note If the test fails to complete successfully, the software analyzes the failure and prints the following error:
External loopback test on interface fc 3/1 failed. Failure reason: Failed to loopback, analysis complete Failed device ID 3 on module 3.

Configuring On-Board Failure Logging

The Generation 2 Fibre Channel switching modules provide the facility to log failure data to persistent storage, which can be retrieved and displayed for analysis. This on-board failure logging (OBFL) feature stores failure and environmental information in nonvolatile memory on the module. The information will help in post-mortem analysis of failed cards.

Verifying System Processes and Logs Configuration

To display the system processes and logs configuration information, perform one of the following tasks:

Command	Purpose
show processes	Displays system processes
show system	Displays system-related status information
show system cores	Display the currently configured scheme for copying cores
show system health	Displays system-related status information
show system health loopback frame-length	Verifies the loopback frequency configuration
show logging onboard status	Displays the configuration status of OBFL

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference*.

This section includes the following topics:

Displaying System Processes

To obtain general information about all processes, follow these steps:

Procedure

Step 1 Choose **Admin > Running Processes**.

You see the Running Processes dialog box.

Where:

- ProcessId = Process ID
- Name = Name of the process
- MemAllocated = Sum of all the dynamically allocated memory that this process has received from the system, including memory that may have been returned
- CPU Time (ms) = CPU time the process has used, in microseconds

Step 2 Click Close to close the dialog box.

Displaying System Status

- In a Cisco MDS 9513 Director, the last four reset-reason codes for the supervisor module in slot 7 and slot 8 are displayed. If either supervisor module is absent, the reset-reason codes for that supervisor module are not displayed.
- In a Cisco MDS 9506 or Cisco MDS 9509 switch, the last four reset-reason codes for the supervisor module in slot 5 and slot 6 are displayed. If either supervisor module is absent, the reset-reason codes for that supervisor module are not displayed.
- In a Cisco MDS 9200 Series switch, the last four reset-reason codes for the supervisor module in slot 1 are displayed.
- The **show system reset-reason module *number*** command displays the last four reset-reason codes for a specific module in a given slot. If a module is absent, then the reset-reason codes for that module are not displayed.
- In a Cisco MDS 9500 Series switch, this command clears the reset-reason information stored in NVRAM in the active and standby supervisor modules.
- In a Cisco MDS 9200 Series switch, this command clears the reset-reason information stored in NVRAM in the active supervisor module.
- Load average—Displays the number of running processes. The average reflects the system load over the past 1, 5, and 15 minutes.
- Processes—Displays the number of processes in the system, and how many are actually running when the command is issued.
- CPU states—Displays the CPU usage percentage in user mode, kernel mode, and idle time in the last one second.
- Memory usage—Displays the total memory, used memory, free memory, memory used for buffers, and memory used for cache in KB. Buffers and cache are also included in the *used* memory statistics.

To display system status from Device Manager, follow these steps:

Procedure

-
- Step 1** Choose **Physical > System**.
You see the System dialog box.
- Step 2** Click Close to close the dialog box.
-

Displaying Core Status

To display cores on a switch, follow these steps:



Note Ensure that SSH2 is enabled on this switch.

Procedure

-
- Step 1** Choose **Admin > Show > Cores**.

You see the Show Cores dialog box.

Module-num shows the slot number on which the core was generated.

Step 2 Click **Close** to close the dialog box.

Additional References

For additional information related to implementing system processes and logs, see the following section:

MIBs


MIBs	MIBs Link
<ul style="list-style-type: none"> CISCO-SYSTEM-EXT-MIB CISCO-SYSTEM-MIB 	<p>To locate and download MIBs, go to the following URL:</p> <p>http://www.cisco.com/en/US/products/ps5989/prod_technical_reference.html</p>

Feature History for System Processes and Logs

[Table 143: Feature History for System Processes and Logs](#), on page 1080 lists the release history for this feature. Only features that were introduced or modified in Release 3.x or a later release appear in the table.

Table 143: Feature History for System Processes and Logs

Feature Name	Releases	Feature Information
Common Information Model	3.3(1a)	Added commands for displaying Common Information Model.
On-line system health maintenance (OHMS) enhancements	3.0(1)	<p>Includes the following OHMS enhancements:</p> <ul style="list-style-type: none"> Configuring the global frame length for loopback test for all modules on the switch. Specifying frame count and frame length on for the loopback test on a specific module. Configuring source and destination ports for external loopback tests. Providing serdes loopback test to check hardware.
On-board failure logging (OBFL)	3.0(1)	Describes OBFL, how to configure it for Generation 2 modules, and how to display the log information.



CHAPTER 57

Configuring QoS

- [Configuring QoS, on page 1081](#)

Configuring QoS

This chapter provides details on the QoS features provided in all switches.

Quality of service (QoS) offers the following advantages:

- Provides relative bandwidth guarantee to application traffic.
- Controls latency experienced by application traffic.
- Prioritizes one application over another (for example, prioritizing transactional traffic over bulk traffic) through bandwidth and latency differentiation.

This chapter includes the following topics:

Information About QoS

QoS implementation in the Cisco MDS 9000 Family follows the differentiated services (DiffServ) model. The DiffServ standard is defined in RFCs 2474 and 2475. The Cisco MDS 9000 Family supports QoS for internally and externally generated control traffic. Within a switch, control traffic is sourced to the supervisor module and is treated as a high priority frame. By default, the QoS feature for certain critical control traffic is enabled. These critical control frames are assigned the highest (absolute) priority. A high priority status provides absolute priority over all other traffic and is assigned in the following cases:

- Internally generated time-critical control traffic (mostly Class F frames).
- Externally generated time-critical control traffic entering a switch in the Cisco MDS 9000 Family from a another vendor's switch. High priority frames originating from other vendor switches are marked as high priority as they enter a switch in the Cisco MDS 9000 Family.

Quality of service (QoS) offers the following advantages:

- Provides relative bandwidth guarantee to application traffic.
- Controls latency experienced by application traffic.

Prioritizes one application over another (for example, prioritizing transactional traffic over bulk traffic) through bandwidth and latency differentiation.

Configuring QoS

This section includes the following topics:

Information About Control Traffic

The Cisco MDS 9000 Family supports QoS for internally and externally generated control traffic. Within a switch, control traffic is sourced to the supervisor module and is treated as a high priority frame. A high priority status provides absolute priority over all other traffic and is assigned in the following cases:

- Internally generated time-critical control traffic (mostly Class F frames).
- Externally generated time-critical control traffic entering a switch in the Cisco MDS 9000 Family from a another vendor's switch. High priority frames originating from other vendor switches are marked as high priority as they enter a switch in the Cisco MDS 9000 Family.

The Cisco MDS 9000 Family supports QoS for internally and externally generated control traffic. Within a switch, control traffic is sourced to the supervisor module and is treated as a high priority frame. A high priority status provides absolute priority over all other traffic and is assigned in the following cases:

- Internally generated time-critical control traffic (mostly Class F frames).
- Externally generated time-critical control traffic entering a switch in the Cisco MDS 9000 Family from a another vendor's switch. High priority frames originating from other vendor switches are marked as high priority as they enter a switch in the Cisco MDS 9000 Family.

Enabling or Disabling Control Traffic

By default, the QoS feature for certain critical control traffic is enabled. These critical control frames are assigned the highest (absolute) priority.



Tip

We do not recommend disabling this feature as all critical control traffic is automatically assigned the lowest priority once you issue this command.

To enable or disable the high priority assignment for control traffic using Fabric Manager, follow these steps:

Procedure

- Step 1** Expand **Switches**, expand **FC Services** and then select **QoS** in the Physical Attributes pane.
The QoS control traffic information is displayed in the Information pane. The **Control** tab is default.
- Step 2** Select the switch on which you want to enable or disable control traffic.
- Step 3** In the Command column, click the drop-down menu and select **enable** or **disable**.
- Step 4** Click **Apply Changes to save your changes**.

Information About Data Traffic

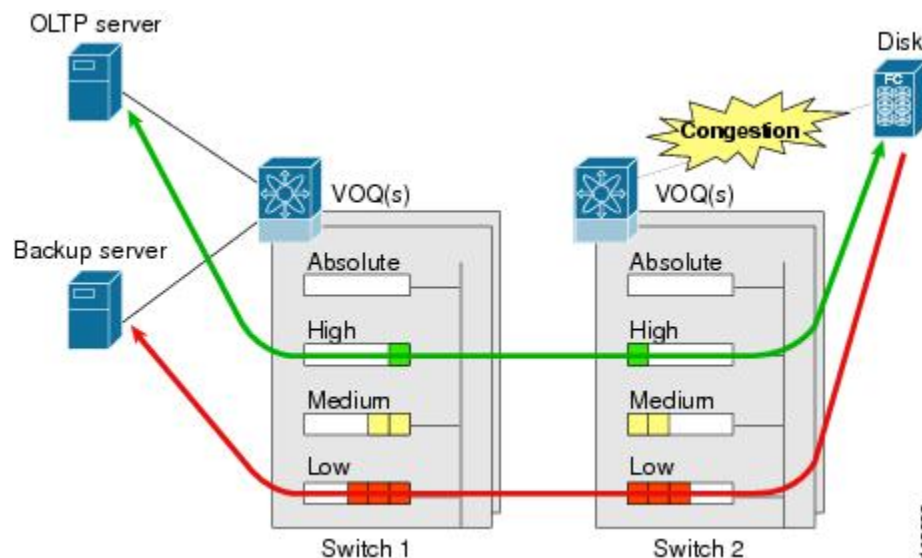
Data traffic can be prioritized in distinct levels of service differentiation: low, medium, or high priority. You can apply QoS to ensure that Fibre Channel data traffic for your latency-sensitive applications receive higher priority over throughput-intensive applications such as data warehousing. With a deficit weighted round robin

(DWRR) scheduler you can ensure that high priority traffic is treated better than low priority traffic. Online transaction processing (OLTP), which is a low volume, latency sensitive application, requires quick access to requested information. For example, DWRR weights of 70:20:10 implies that the high priority queue is serviced at 7 times the rate of the low priority queue. This guarantees lower delays and higher bandwidths to high priority traffic if congestion sets in. A similar configuration in the second switch ensures the same traffic treatment in the other direction.

Online transaction processing (OLTP), which is a low volume, latency sensitive application, requires quick access to requested information. Backup processing application require high bandwidth but are not sensitive to latency. In a network that does not support service differentiation, all traffic is treated identically—they experience similar latency and are allocated similar bandwidths. The QoS feature in the Cisco MDS 9000 Family switches provides these guarantees.

Data traffic can be prioritized in distinct levels of service differentiation: low, medium, or high priority. You can apply QoS to ensure that Fibre Channel data traffic for your latency-sensitive applications receive higher priority over throughput-intensive applications such as data warehousing.

Figure 184: Prioritizing Data Traffic



In this image, the OLTP traffic arriving at Switch 1 is marked with a high priority level of throughput classification (class map) and marking (policy map). Similarly, the backup traffic is marked with a low priority level. The traffic is sent to the corresponding priority queue within a virtual output queue (VOQ).

A deficit weighted round robin (DWRR) scheduler configured in the first switch ensures that high priority traffic is treated better than low priority traffic. For example, DWRR weights of 70:20:10 implies that the high priority queue is serviced at 7 times the rate of the low priority queue. This guarantees lower delays and higher bandwidths to high priority traffic if congestion sets in. A similar configuration in the second switch ensures the same traffic treatment in the other direction.

If the ISL is congested when the OLTP server sends a request, the request is queued in the high priority queue and is serviced almost immediately since the high priority queue is not congested. The scheduler assigns its priority over the backup traffic in the low priority queue.



Note When the high priority queue does not have traffic flowing through, the low priority queue uses all the bandwidth and is not restricted to the configured value.

A similar occurrence in Switch 2 sends a response to the transaction request. The round trip delay experienced by the OLTP server is independent of the volume of low priority traffic or the ISL congestion. The backup traffic uses the available ISL bandwidth when it is not used by the OLTP traffic.

Comparing VSAN Versus Zone-Based QoS

While you can configure both zone-based QoS and VSAN-based QoS configurations in the same switch, both configurations have significant differences. The following table highlights the differences between configuring QoS priorities based on VSANs versus zones.

Table 144: QoS Configuration Differences

VSAN-Based QoS	Zone-Based QoS
If you configure the active zone set on a given VSAN and also configure QoS parameters in any of the member zones, you cannot associate the policy map with the VSAN.	You cannot activate a zone set on a VSAN that already has a policy map associated.
If the same flow is present in two class maps associated to a policy map, the QoS value of the class map attached first takes effect.	If the same flow is present in two zones in a given zone set with different QoS values, the higher QoS value is considered.
—	During a zone merge, if the Cisco NX-OS software detects a mismatch for the QoS parameter, the link is isolated.
Takes effect only when QoS is enabled.	Takes effect only when QoS is enabled.

Configuring Data Traffic

To configure QoS using Fabric Manager, follow these steps:

Procedure

-
- Step 1** Enable the QoS feature.
 - Step 2** Create and define class maps.
 - Step 3** Define service policies.
 - Step 4** Apply the configuration.
-

Information About Class Map Creation

Using the class map feature you can create and define traffic class with match criteria to identify traffic belonging to that class. The class map name is restricted to 63 alphanumeric characters and defaults to the match-all option. Flow-based traffic uses one of the following values:

- WWN—The source WWN or the destination WWN.
- Fibre Channel ID (FC ID) —The source ID (SID) or the destination ID (DID).
- Source interface—The ingress interface.

Use the class map feature to create and define a traffic class with match criteria to identify traffic belonging to that class. The class map name is restricted to 63 alphanumeric characters and defaults to the match-all option. Flow-based traffic uses one of the following values:

- WWN—The source WWN or the destination WWN.
- Fibre Channel ID (FC ID) —The source ID (SID) or the destination ID (DID). The possible values for mask are FFFFFFFF (the entire FC ID is used—this is the default), FFFF00 (only domain and area FC ID is used), or FF0000 (only domain FC ID is used).



Note An SID or DID of 0x000000 is not allowed.

- Source interface—The ingress interface.



Tip The order of entries to be matched within a class map is not significant.

Creating a Class Map

To create a class map, follow these steps:

Procedure

- Step 1** Expand **Switches**, expand **FC Services** and then select **QoS** in the Physical Attributes pane. The QoS information is displayed in the Information pane. The **Control** tab is the default.
- Step 2** In the **Class Maps** tab, click **Create Row** to create a new class map. You see the Create Class Maps dialog box.
- Step 3** Select the switches for the class map.
- Step 4** Enter the source ID or the destination **ID** in the field.
- Step 5** Enter a name for the class map.
- Step 6** Select a Match mode. You can either match **any** or **all** criterion with one match statement from the class map configuration mode.
- Step 7** Click **Create** to proceed with creating the class map.

Information About Service Policy Definition

Service policies are specified using policy maps. Policy maps provide an ordered mapping of class maps to service levels. The order of the class maps within a policy map is important to determine the order in which the frame is compared to class maps. The first matching class map has the corresponding priority marked in the frame. You can specify multiple class maps within a policy map, and map a class map to a high, medium, or low service level. The default priority is low. Alternatively, you can map a class map to a differentiated services code point (DSCP). The DSCP is an indicator of the service level for a specified frame.

Service policies are specified using policy maps. Policy maps provide an ordered mapping of class maps to service levels. You can specify multiple class maps within a policy map, and map a class map to a high, medium, or low service level. The default priority is low. The policy map name is restricted to 63 alphanumeric characters.

As an alternative, you can map a class map to a differentiated services code point (DSCP). The DSCP is an indicator of the service level for a specified frame. The DSCP value ranges from 0 to 63, and the default is 0. A DSCP value of 46 is disallowed.

The order of the class maps within a policy map is important to determine the order in which the frame is compared to class maps. The first matching class map has the corresponding priority marked in the frame.



Note Refer to Implementing [Quality of Service Policies with DSCP](#) for further information on implementing QoS DSCP values.



Note Class maps are processed in the order in which they are configured in each policy map.

About Service Policy Enforcement

When you have configured a QoS data traffic policy, you must enforce the data traffic configuration by applying that policy to the required VSAN(s). If you do not apply the policy to a VSAN, the data traffic configuration is not enforced. You can only apply one policy map to a VSAN.



Note You can apply the same policy to a range of VSANs.

About the DWRR Traffic Scheduler Queue

The DWRR scheduler services the queues in the ratio of the configured weights. Higher weights translate to proportionally higher bandwidth and lower latency. The default weights are 50 for the high queue, 30 for the medium queue, and 20 for the low queue. The Cisco NX-OS software supports four scheduling queues:

- Strict priority queues are queues that are serviced in preference to other queues—it is always serviced if there is a frame queued in it regardless of the state of the other queues.
- QoS assigns all other traffic to the DWRR scheduling high, medium, and low priority traffic queues.

The Cisco NX-OS software supports four scheduling queues:

- Strict priority queues are queues that are serviced in preference to other queues—it is always serviced if there is a frame queued in it regardless of the state of the other queues.
- QoS assigns all other traffic to the DWRR scheduling high, medium, and low priority traffic queues.

The DWRR scheduler services the queues in the ratio of the configured weights. Higher weights translate to proportionally higher bandwidth and lower latency. The default weights are 50 for the high queue, 30 for the medium queue, and 20 for the low queue. Decreasing order of queue weights is mandated to ensure the higher priority queues have a higher service level, though the ratio of the configured weights can vary (for example, one can configure 70:30:5 or 60:50:10 but not 50:70:10).

The following table describes the QoS behavior for Generation 1, Generation 2, and Generation 3 switching modules.

Table 145: QoS Behavior for Generation 1 and Generation 2 Switching Modules

Source Module Type	Destination Module Type	QoS Behavior Description
Generation 1	Generation 1	QoS behavior reflects the DWRR configuration for traffic coming in through a given port and queued to the same egress port. All the other traffic share equal bandwidth.
Generation 1	Generation 2 or Generation 3	QoS behavior reflects the DWRR configuration for traffic coming in through a given port and queued to the same egress port. All the other streams share equal bandwidth.
Generation 2 or Generation 3	Generation 1	Bandwidth partitioning is equal for all the traffic.
Generation 2 or Generation 3	Generation 2 or Generation 3	QoS behavior reflects the DWRR weights configuration for all possible streams.

Changing the Weight in a DWRR Queue

To change the weight in a DWRR queue using Fabric Manager, follow these steps:

Procedure

-
- Step 1** Expand **Switches**, expand **FC Services** and then select **QoS** in the Physical Attributes pane.
The QoS control traffic information is displayed in the Information pane. The default is the **Control** tab.
- Step 2** Click the **DWRR** tab.
You see the queue status and weight.
- Step 3** Select a switch and change the weight.
- Step 4** Click the **Apply Changes** icon to save your changes.
-

Limiting Ingress Port Rate Limiting

About Limiting Ingress Port Rate

A port rate limiting feature helps control the bandwidth for individual Fibre Channel ports. Port rate limiting is also referred to as ingress rate limiting because it controls ingress traffic into a Fibre Channel port. The feature controls traffic flow by limiting the number of frames that are transmitted out of the exit point on the MAC. Port rate limiting works on all Fibre Channel ports. The rate limit ranges from 1 to 100% and the default is 100%.

A port rate limiting feature helps control the bandwidth for individual Fibre Channel ports. Port rate limiting is also referred to as ingress rate limiting because it controls ingress traffic into a Fibre Channel port. The feature controls traffic flow by limiting the number of frames that are transmitted out of the exit point on the MAC. Port rate limiting works on all Fibre Channel ports. The rate limit ranges from 1 to 100% and the default is 100%.



Note


Port rate limiting can only be configured on Cisco MDS 9100 Series switches, Cisco MDS 9216i switches, and MPS-14/2 modules.

This feature can only be configured if the QoS feature is enabled and if this configuration is performed on a Cisco MDS 9100 series switch, Cisco MDS 9216i switch, or MPS-14/2 module.

To configure the port rate limiting value using Fabric Manager, follow these steps:

Procedure

- Step 1** Expand **Switches**, expand **FC Services** and then select **QoS** in the Physical Attributes pane.
The QoS control traffic information is displayed in the Information pane. The default is the **Control** tab.
- Step 2** Click the **Rate Limit** tab.
- Step 3** Select the switch whose port rate limit you want to change.
- Step 4** Enter the desired port rate limit in the Percent column.
- Step 5** Click the **Apply Changes** icon to save your changes.



CHAPTER 58

Configuring Port Tracking

- [Configuring Port Tracking, on page 1089](#)

Configuring Port Tracking

The port tracking feature is unique to the Cisco MDS 9000 Family of switches. This feature uses information about the operational state of the link to initiate a failure in the link that connects the edge device. This process of converting the indirect failure to a direct failure triggers a faster recovery process toward 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

The Port Tracking feature is unique to the Cisco MDS 9000 Family of switches. This feature uses information about the operational state of the link to initiate a failure in the link that connects the edge device. This process of converting the indirect failure to a direct failure triggers a faster recovery process towards redundant links. Generally, hosts can instantly recover from a link failure on a link that is immediately (direct link) connected to a switch. However, recovering from an indirect link failure between switches in a WAN or MAN fabric with a keep-alive mechanism is dependent on several factors such as the time out values (TOVs) and on registered state change notification (RSCN) information. 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.

The port tracking feature is disabled by default in all switches in the Cisco 9000 Family. When you enable this feature, port tracking is globally enabled for the entire switch.

Generally, hosts can instantly recover from a link failure on a link that is immediately (direct link) connected to a switch. However, recovering from an indirect link failure between switches in a WAN or MAN fabric with a keep-alive mechanism is dependent on several factors such as the time out values (TOVs) and on registered state change notification (RSCN) information.

In [Figure 185: Traffic Recovery Using Port Tracking, on page 1090](#), when the direct link 1 to the host fails, recovery can be immediate. However, when the ISL 2 fails between the two switches, recovery depends on TOVs, RSCNs, and other factors.

Figure 185: Traffic Recovery Using Port Tracking

The port tracking feature monitors and detects failures that cause topology changes and brings down the links connecting the attached devices. When you enable this feature and explicitly configure the linked and tracked ports, the Cisco NX-OS software monitors the tracked ports and alters the operational state of the linked ports on detecting a link state change.

The following terms are used in this chapter:

- **Tracked ports:** A port whose operational state is continuously monitored. The operational state of the tracked port is used to alter the operational state of one or more ports. Fibre Channel, VSAN, PortChannel, FCIP, or a Gigabit Ethernet port can be tracked. Generally, ports in E and TE port modes can also be Fx ports.
- **Linked ports:** A port whose operational state is altered based on the operational state of the tracked ports. Only a Fibre Channel port can be linked.

Guidelines and Limitations

Before configuring port tracking, consider the following guidelines:

- Verify that the tracked ports and the linked ports are on the same Cisco MDS switch.
- Do not track a linked port back to itself (for example, Port fc1/2 to Port fc2/5 and back to Port fc1/2) to avoid recursive dependency.
- Be aware that the linked port is automatically brought down when the tracked port goes down. Be aware that the linked port is automatically brought down when the tracked port goes down.

Default Settings

[Table 146: Default Port Tracking Parameters](#), on page 1090 lists the default settings for port tracking parameters.

Table 146: Default Port Tracking Parameters

Parameters	Default
Port tracking	Disabled.
Operational binding	Enabled along with port tracking.

Configuring Port Tracking

Port tracking has the following features:

- The application brings the linked port down when the tracked port goes down. When the tracked port recovers from the failure and comes back up again, the tracked port is also brought up automatically (unless otherwise configured).
- You can forcefully continue to keep the linked port down, even though the tracked port comes back up. In this case, you must explicitly bring the port up when required.

This section includes the following topics:

Enabling Port Tracking

The port tracking feature is disabled by default in all switches in the Cisco 9000 Family. When you enable this feature, port tracking is globally enabled for the entire switch.

To configure port tracking, enable the port tracking feature and configure the linked port(s) for the tracked port.

To enable port tracking with Fabric Manager, follow these steps:

Procedure

Step 1 Expand **Switches**, expand **Interfaces**, and then select **Port Tracking** in the Physical Attributes pane.

The port tracking information is displayed in the Information pane shown in [Figure 186: Port Tracking, on page 1091](#). The default is the **Controls** tab.

Figure 186: 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 Click in the Command column to **enable** or **disable** port tracking.

Depending on your selection the corresponding entry in the Status column changes.

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

The entry in the Result column changes to **success**.

Information About Configuring Linked Ports

You can link ports using one of two methods:

- Operationally binding the linked port(s) to the tracked port (default).

- Continuing to keep the linked port down forcefully—even if the tracked port has recovered from the link failure.

Information About Tracked Port

A tracked port is one whose operational state is continuously monitored. The operational state of the tracked port is used to alter the operational state of one or more ports. Fibre Channel, VSAN, PortChannel, FCIP, or a Gigabit Ethernet port can be tracked. Generally, ports in E and TE port modes can also be Fx ports. Ports can be linked using one of the following two methods:

- Operationally binding the linked port(s) to the tracked port (default). Operational binding comes to effect automatically when you configure the first tracked port.
- Continuing to keep the linked port down forcefully—even if the tracked port has recovered from the link failure.

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, follow these steps:

Procedure

-
- | | |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Expand Switches , expand Interfaces , and then select Port Tracking in the Physical Attributes pane. The port tracking information is displayed in the Information pane. The default is the Controls tab. |
| Step 2 | Click the Dependencies tab. |
| Step 3 | Click Create Row .
You see the Create Port Tracking Dependencies dialog box. |
| Step 4 | Select the switch whose ports you want to track by and selecting a switch from the drop-down list. |
| Step 5 | Select the linked port(s) that should be bound to the tracked port(s) by clicking the browse button and selecting 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 Information About Monitoring Ports in a VSAN, on page 1093 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 Forceshutdown check box if you want to forcefully shutdown the tracked port. |
| Step 9 | Click Create to proceed with creating this dependency.
If tracking is established, you see Success in the lower left corner of the dialog box. |
| Step 10 | Click Close to close the dialog box. |
-

Information 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.

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 the following image, 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.

Information 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.

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.

Information 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.

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 PortDetailed Steps

To forcefully shut down a tracked port, see [Information About Tracked Port, on page 1092](#).



CHAPTER 59

Configuring FlexAttach Virtual pWWN

- [Configuring FlexAttach Virtual pWWN, on page 1095](#)

Configuring FlexAttach Virtual pWWN

This chapter describes how to configure the FlexAttach virtual port world-wide name (pWWN) feature.

Information About FlexAttach Virtual pWWN

This section includes the following topics:

FlexAttach Virtual pWWN

FlexAttach virtual pWWN feature facilitates server and configuration management. In a SAN environment, the server installation or replacement, requires interaction and coordination among the SAN and server administrators. For coordination, it is important that the SAN configuration does not change when a new server is installed, or when an existing server is replaced. FlexAttach virtual pWWN minimizes the interaction between the server administrator and the SAN administrator by abstracting the real pWWN using virtual pWWNs.

When FlexAttach virtual pWWN is enabled on an interface, a virtual pWWN is assigned to the server interface. The real pWWN is replaced by a virtual pWWN, which is used for a SAN configuration such as zoning.

Server administrators can benefit from FlexAttach in the following scenarios:

- **Pre-configure**—Pre-configure SAN for new servers that are not available physically yet. For example, they may be on order. FlexAttach can be enabled on the ports designated for the new servers and use the virtual WWNs assigned for configuring SAN. The new servers are then plugged into the fabric without any change needed in the SAN.
- **Replacement to the same port**—A failed server can be replaced onto the same port without changing the SAN. The new server gets a same pWWN as the failed server because the virtual pWWN is assigned to the port.
- **Replacement to (spare)**—A spare server, which is on the same NPV device or a different NPV device) can be brought online without changes to the SAN. This action is achieved by moving the virtual port WWN from the current server port to the spare port.
- **Server Mobility**—A server can be moved to another port on the same NPV device or another NPV device without changing the SAN. This is accomplished by moving the virtual pWWN to the new port. No

change is needed if FlexAttach was configured using the physical port WWN of the server to the virtual port WWN mapping.

Difference Between San Device Virtualization and FlexAttach Port Virtualization

The following table describes the difference between SAN device virtualization (SDV) and FlexAttach port virtualization.

Table 147: Difference Between SDV and FlexAttach Virtualization

SAN Device Virtualization (SDV)	FlexAttach Virtualization
Facilitates target and disk management, and only facilitates disk and data migration.	Facilitates server management and has no restriction on the end devices used.
WWN NAT and Fibre Channel ID (FC-ID) are allocated on the virtual device, both primary and secondary.	WWN and Network Address Transport (NAT) is allocated to host bus adapter (HBA).
FC-ID rewrite on the switch indicates a rewrite-capable switch on the path.	No rewrite requirements.
Configuration is distributed. This allows programming rewrites and connectivity anywhere.	Configuration distribution is not required for any of the interface-based configurations.
Configuration is secured to device alias.	Does not require device alias for virtual pWWN.
Does not allow automapping to the secondary device.	Allows automapping to the new HBA. Mapping process is manual for NPIV.

FlexAttach Virtual pWWN CFS Distribution

The FlexAttach virtual pWWN configuration is distributed for CFS through IPv4, and is enabled by default. The FlexAttach virtual pWWN distribution, by default, is on CFS region 201. The CFS region 201 links only to the NPV-enabled switches. Other CFS features such as syslog is on region 0. Region 0 will be linked through IPv4 for all NPV switches on the same physical fabric. If CFS has an option to link through IPv4 or ISL, then CFS will select the ISL path.



Note

NPV switches do not have ISL (E or TE ports) and are linked through IPv4.

Security Settings for FlexAttach Virtual pWWN

Security settings for the FlexAttach virtual pWWN feature are done by port security at the NPV core. Node WWN of the end device is used to provide physical security.

For more details on enabling port security, refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.

Guidelines and Limitations

Following are recommended guidelines and requirements when deploying FlexAttach virtual pWWN:

- FlexAttach configuration is supported only on NPV switches.

- Cisco Fabric Services (CFS) IP version 4 (IPv4) distribution should be enabled.
- Virtual WWNs should be unique across the fabric.

Configuring FlexAttach Virtual pWWN

This section includes the following topics:

Automatically Assigning FlexAttach Virtual pWWN

Automatic assignment of virtual pWWN can be configured on an NPV switch globally, per VSAN, or per port. When assigned automatically, a virtual WWN is generated from the device local switch WWN.

Before you begin

The port must be in a shut state when the virtual pWWN is enabled.

To enable virtual pWWN automatically for all the interfaces, follow these steps:

Procedure

-
- | | |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | From the Device Manger menu bar, select FC > FlexAttach .
You see the FlexAttach window. |
| Step 2 | Check the VirtualWwnAuto check box to enable automatic generation of virtual WWNs on all the fabric port interfaces. |
| Step 3 | troubleshooting: <ul style="list-style-type: none">• When the <i>interface-list</i> value is not included in the command, virtual pWWN is assigned globally.• All the interfaces mentioned in the <i>interface-list</i> value must be in a shut state. |
-

Launching FlexAttach in DCNM-SAN

To launch FlexAttach in DCNM-SAN, follow these steps:

Procedure

-
- | | |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | In the Logical Domains pane, select a switch. |
| Step 2 | In the Physical Attributes pane, expand Switches > FC Services > N_Port Virtualizer (NPV) > Flex Attach .
The FlexAttach configuration pane appears to the right. |
-

Manually Assigning FlexAttach Virtual pWWN

You can manually assign a WWN to the interface, without generating it through the switch. Several checks are done by the NPV core to ensure the uniqueness of virtual pWWNs in the switch. When duplicate virtual pWWNs are configured, the subsequent logins are rejected by the NPV core switch.

Before you begin

- Some ports may be in automode, some in manual mode, and the virtual pWWNs need not be assigned.
- The port must be in a shut state when a virtual pWWN is assigned.



Note

- The interface mentioned in the interface value must be in a shut state.

To enable virtual pWWN on each interface manually, follow these steps:

Procedure

- Step 1** In the Physical Attributes pane, expand **Switches > FC Services > N_Port Virtualizer (NPV) > Flex Attach**. The FlexAttach configuration pane appears to the right.
- Step 2** Click the **Virtual pWWNs** tab. The Virtual pWWN tab view displays a list of the interfaces.
- Step 3** Check the **Auto** check box to automatically generate the virtual pWWN value for the selected interface. The virtual port WWN value for the selected interface in DCNM-SAN is automatically generated.

Mapping pWWN to Virtual pWWN

You can configure virtual pWWNs through real pWWNs. This process is required for NPIV hosts containing multiple pWWNs, of which only FLOGI is mapped to the virtual pWWN. Subsequent FDSIDs will have different mappings.

Several checks are done by the NPV core to ensure the uniqueness of virtual pWWNs in the switch across the NPV switches. When duplicate virtual pWWNs are configured, the subsequent logins are rejected by the NPV core switch.

Before you begin

The interface must be in a shut state and the specified virtual pWWN should not be logged in.



Note

The specified virtual pWWN and the real pWWN must not be logged in.

To map pWWN to virtual pWWN, follow these steps:

Procedure

-
- Step 1** In the Physical Attributes pane, expand **Switches > FC Services > N_Port Virtualizer (NPV) > Flex Attach**.
- Step 2** In the FlexAttach window, click the **Physical to Virtual pWWNs** tab.
- You see the **Physical to Virtual pWWNs** tab view.
- The LastChange field displays the time when the virtual pWWN was changed.
-

Using the Server Admin FlexAttach Wizards

As in DCNM for SAN Release 4.1(1) and later, the Server Admin perspective view of the DCNM-SAN GUI provides the following FlexAttach wizards, which the DCNM-SAN users with server-admin role can use to configure FlexAttach:

To access the FlexAttach wizards, follow these steps:

Procedure

-
- Step 1** Log in to DCNM-SAN with a username and password that has the server-admin role assigned.
- Step 2** Discover and open the fabric on which you want to configure FlexAttach.
- Step 3** In the DCNM-SAN window displayed, select **Tools > NPV** to display the list of NPV wizards that includes the Flex Attach options like **Flex Attach Pre-Configure Server**, **Flex Attach Move Server**, and **Flex Attach Replace Server**.
-

Pre-Configuring FlexAttach for a New Server

Using the Pre-configure Server wizard, you can configure FlexAttach for servers that are not physically available currently. FlexAttach can be enabled on the ports designated for the new servers and can use the virtual WWNs assigned for configuring SAN. When the new servers are available, the servers can then be plugged into the fabric without any change needed in the SAN.

The Pre-Configure Server wizard can be used to accomplish the following tasks:

Pre-Configuring FlexAttach for All the Ports

Using the Pre-Configure Server Basic configuration wizard, you can set the following port configurations for all the ports in one or more switches in common:

- Enable or disable FlexAttach Auto on all ports.
- Set the default VSAN ID for all the ports.
- Set the interface status for all the ports.

To pre-configure a common setting for all the ports in one or more switches, follow these steps:

Procedure

-
- Step 1** In the DCNM-SAN window, select **Tools > NPV > Flex Attach Pre-configure Server**.
The Pre-Configure Wizard is displayed.
- Step 2** In the Pre-Configure Server window, click the **Basic** radio button to configure a common setting to all the ports on one or more switches. Click **Next**.
The Basic Configuration window is displayed.
- Step 3** In the Basic Configuration window, check the check box to select one or more switches from the list of NPV switches in the fabric.
- Step 4** Check the **Enable FlexAttach Auto on every port** check box to enable FlexAttach on all the ports of all the selected switches.
- Step 5** (Optional) From the VSAN ID drop-down list, select a VSAN ID to assign the selected VSAN ID to all the ports.
- Note** Only the set of VSANs to which all the selected switches belong are listed. If no VSAN ID is selected, then the existing VSAN configuration is retained.
- Step 6** Click the **Up** or **Down** radio button to assign the selected interface status.
- Note** The status of only the F ports in the selected switches will be brought to up or down state.
- Step 7** Click **Finish** to pre-configure the selected settings to all the ports on all the selected switches.
The Configuration window is displayed with the finished applying the new configuration message.
-

Pre-Configuring FlexAttach for Each Port Individually

Using the Pre-Configure Server Advanced configuration wizard, you can set the following port configurations for each port in one or more switches individually:

- Enable FlexAttach Auto on all ports.
- Enable FlexAttach Auto or Manual on individual ports.
- Set the virtual pWWN for ports where FlexAttach is enabled Manually.
- Set pWWN to vPWWN mapping.
- Set the default VSAN ID for each port.
- Set the Interface status for each port.

To pre-configure FlexAttach on each port individually, follow these steps:

Procedure

-
- Step 1** In the DCNM-SAN window, select **Tools > NPV > Flex Attach Pre-configure Server**.
The Pre-Configure Server window is displayed.
- Step 2** In the Pre-Configure Server window, click the **Advanced** radio button to configure FlexAttach on each port individually. Click **Next**.

The Pre-Configure Server Advanced configuration window is displayed.

Note From the **Interface tab**, you can select a switch from the list of switches displayed in the left pane and click **Disable FlexAttach** to change the switches to manual configuration. Select **Undo Changes** to return to the previous configuration.

Step 3 In the **Interface tab**, and select a switch from the list of switches displayed in the left pane.
The switch configuration details are displayed in the right pane with tabs and columns.

Step 4 Configure the following settings, for each interface:

- In the Status Admin column corresponding to the interface, double-click and then select up or down from the drop-down list.
- In the VSAN column corresponding to the interface, double-click and then select the VSAN ID from the drop-down list of existing VSAN IDs.
- In the Auto column corresponding to the interface, double-click and then select Auto to automatically enable FlexAttach or select Manual to manually enable FlexAttach later.
- In the Interface vPWWN cell, enter the vPWWN if Manual was selected in the Auto FlexAttach configuration cell.

Note You can click **Set All Auto** to change all the interfaces with manual FlexAttach configuration to Auto on the selected switch. However, if a valid virtual PWWN value is already configured, then changing it to Auto does not change the configuration. Before you change from Manual to Auto, update the Interface vPWWN column with the 00:00:00:00:00:00:00:00 value.

Step 5 Repeat Step 3 through Step 4 for each switch.

Step 6 Click the **pWWN to vPWWN** tab to configure pWWN to virtual PWWN mapping.

The Advanced Configuration window is displayed.

Step 7 From the Select a switch drop-down list, select the switch to display the existing pWWN to Virtual PWWN mapping table for the CFS region to which the switch belongs, and then follow these steps to add virtual pWWN to virtual pWWN automap entries:

- Click **Add Row** to display the pWWN to vPWWN dialog box.
- Enter the pWWN and the corresponding Virtual PWWN.
- Click **Create** to add the mapping list.

Note To delete an existing mapping, select the row, and then click **Delete Row**. Only one pWWN to vPWWN table can be updated at a time. To update the table for each CFS region, perform Step 6 through Step 8 for a switch from each CFS region.

Step 8 Click **Finish** to complete the configurations for each port.

Moving a Server to Another Port or Switch

Using the Move Server wizard, you can move a server to another port on the same NPV device or another NPV device without changing the SAN. This is accomplished by moving the virtual pWWN to the new port. No change is needed if FlexAttach was configured using the physical port WWN of the server to the virtual port WWN mapping.

To move a server to a different port in the same switch, or in a different switch, follow these steps:

Procedure

-
- Step 1** In the DCNM-SAN window, select **Tools > NPV > FlexAttach Move Server**.
The Move Server wizard is displayed.
- Step 2** In the Move Server to.. window, click the **Another Port on the Same Switch** radio button or click the **Another Port on a Different Switch** radio button.
- Step 3** Click **Next**.
The Choose Old Port window is displayed.
- Step 4** From the Select a switch drop-down list, select a switch.
The switch ports are listed. To support moving a server from a failed port that is in down state, the ports in down state are also listed.
- Step 5** From the list of interfaces, select the port from which you want to move the server from.
- Step 6** Click **Next**.
The Choose New Port window is displayed.
- Step 7** From the Select a switch drop-down list box, select a switch.
Note If the **Another Port on the Same Switch** radio button was chosen, then the Select a switch drop-down list is disabled.
- Step 8** From the list of interfaces, select the port to which you want to move the server to.
- Step 9** Click **Next**.
The Server WWN window is displayed.
In the Server WWN window, if the FlexAttach global mapping table is empty, the wizard automatically prefills the drop-down table with the interface virtual PWWN of the source port, and the virtual PWWN field is not editable.
If the FlexAttach global mapping table is not empty, the VPWWN field is blank and editable. From the drop-down list that displays all existing entries from the global mapping table, select the VPWWN entry or type the required entry.
- Step 10** Click **Finish**.
-

Replacing a Server with Another Server

You can use the Replace Server wizard to accomplish the following tasks:

- Replace a failed server with a new server onto the same port without changing the SAN. The new server gets the same virtual pWWN as the failed server because the virtual pWWN is assigned to the port.
- Replace a server with a spare server on the same NPV device or a different NPV device, which can be brought online without changes to the SAN. This is achieved by moving the virtual port WWN from the current server port to the spare port.

Replacing a Server on the Same Port

To replace a failed server with a new server on the same port, follow these steps:

Procedure

- Step 1** In the DCNM-SAN window, select **Tools > NPV > FlexAttach Replace Server**.
The Replace Failed Server window is displayed.
- Step 2** In the Replace Server Wizard, click the **On Same Port** radio button.
- Step 3** Click **Next**.
The Choose Failed Port window is displayed.
- Step 4** In the Choose Failed Port selection window, from the Select a switch drop-down list, select the switch.
- Step 5** From the list of interfaces displayed, select the port on which the server needs to be replaced.
- Step 6** Click **Next**.
The Server WWN window is displayed.
- Step 7** In the Server WWN window, enter the existing FlexAttach server virtual port WWN that needs to be replaced, and the new server physical port WWN.
- Step 8** Click **Finish** to complete the FlexAttach configuration for the new server.
-

Replacing the Server to a Different Port on the Same Switch

To replace a server with a spare server on a different port in the same switch, follow these steps:

Procedure

- Step 1** In the DCNM-SAN window, select **Tools > NPV > name=">FlexAttach name=">Replace Server**.
The Replace Failed Server wizard is displayed.
- Step 2** In the Replace Failed Server wizard, click the **With Spare Server on Same NPV Switch** radio button.
- Step 3** Click **Next**.
The Choose Failed Port window is displayed.
- Step 4** In the Choose Failed Port selection window, from the Select a Switch drop-down list, select the switch.
- Step 5** From the list of interfaces displayed, select the port from which the server needs to be detached.
- Step 6** Click **Next**.
The Choose New Port window is displayed.
- Step 7** In the Choose New Port selection window, select the port on which the spare server is connected.
- Step 8** Click **Next**.
The Server WWN window is displayed.

In the Server WWN window, if the FlexAttach global mapping table is empty, the wizard automatically prefills the drop-down table with the interface virtual vPWWN of the source port to be replaced, and the vPWWN field is not editable. In this case, the **Allow wizard to change from “pWWN to vPWWN” mapping to “interface to vPWWN” mapping** is treated as true.

If the FlexAttach global mapping table is not empty, the vPWWN field is blank and editable. From the drop-down list that displays all existing entries from the global mapping table, select the vPWWN entry or type the required entry, and the new server physical port WWN.

Check the **Allow wizard to change from “pWWN to vPWWN” mapping to “interface to vPWWN” mapping** check box to remove the pWWN to vPWWN entry from the CFS Region mapping table, and configure the mapping only at the interface.

- Step 9** Click **Finish** to complete the FlexAttach configuration for the spare server.

Replacing with a Server on a Different Switch

To replace a server with a spare server on a different switch, follow these steps:

Procedure

- Step 1** In the DCNM-SAN window, select **Tools > NPV > FlexAttach Replace Server**.
The Replace Server wizard is displayed.
- Step 2** In the Replace Server wizard, click the **With Spare Server on a Different NPV switch** radio button.
- Step 3** Click **Next**.
The Failed Server Port window is displayed.
- Step 4** In the Failed Server Port selection window, from the Select a Switch drop-down list, select the switch.
- Step 5** From the list of interfaces displayed, select the port from which the server needs to be detached.
- Step 6** Click **Next**.
The New Port window is displayed.
- Step 7** In the New Port selection window, select the switch and the port on which the spare server is connected.
- Step 8** Click **Next**.
The Server WWN window is displayed.
- In the Server WWN window, if the FlexAttach global mapping table is empty, the wizard automatically prefills the table with the interface virtual vPWWN of the source port to be replaced, and the vPWWN field is not editable. In this case, the **Allow wizard to change from “pWWN to vPWWN” mapping to “interface to vPWWN” mapping** is treated as true.
- If the FlexAttach global mapping table is not empty, the vPWWN field is blank and editable. From the drop-down list box which displays all existing entries from the global mapping table, select the vPWWN entry or type the required entry, and the new server physical port WWN.
- Check the **Allow wizard to change from “pWWN to vPWWN” mapping to “interface to vPWWN” mapping** check box to remove the pWWN to vPWWN entry from the CFS Region mapping table, and configure the mapping only at the interface.

Step 9 Click **Finish** to complete the FlexAttach configuration for the spare server.



CHAPTER 60

Configuring Interface Buffers

- [Configuring Interface Buffers, on page 1107](#)

Configuring Interface Buffers

This chapter describes how to configure buffer credits for the Fibre Channel interfaces.

This chapter includes the following topics:

Information About Interface Buffers

Fibre Channel interfaces use buffer credits to ensure all packets are delivered to their destination.

This section includes the following topics:

Buffer-to-Buffer Credits

Buffer-to-buffer credits (BB_credits) are a flow-control mechanism to ensure that Fibre Channel switches do not run out of buffers, so that switches do not drop frames. BB_credits are negotiated on a per-hop basis.

The receive BB_credit (fcrxbbcredit) value may be configured for each Fibre Channel interface. In most cases, you do not need to modify the default configuration.

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

- For 16-port switching modules and full rate ports, the default value is 16 for Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required.
- For 32-port switching modules and host-optimized ports, the default value is 12 for Fx, E, and TE modes. These values cannot be changed.
- For Generation 2, Generation 3, and Generation 4 switching modules, see the [Buffer Pools, on page 1108](#).



Note

In the Cisco MDS 9100 Series switches, the groups of ports on the left outlined in white are in dedicated rate mode. The other ports are host-optimized. Each group of 4 host-optimized ports have the same features as for the 32-port switching module.



Note Because Generation 1 modules do not support as many buffer-to-buffer credits as Generation 4 modules supports, you cannot configure an ISL on E or TE ports between a Generation 1 module such as the 16-port 1-, 2-Gbps Fibre Channel Switching Module (DS-X9016) and a Generation 4 module such as the 48 port 8-Gbps Advanced Fibre Channel module (DS-X9248-256K9) or the 32-port 8-Gbps Advanced Fibre Channel module (DS-X9232-256K9).

Performance Buffers

Regardless of the configured receive BB_credit value, additional buffers, called performance buffers, improve switch port performance. Instead of relying on the built-in switch algorithm, you can manually configure the performance buffer value for specific applications (for example, forwarding frames over FCIP interfaces).



Note Performance buffers are not supported on the Cisco MDS 9148 Fabric Switch, Cisco MDS 9124 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

For each physical Fibre Channel interface in any switch in the Cisco MDS 9000 Family, you can specify the amount of performance buffers allocated in addition to the configured receive BB_credit value.

The default performance buffer value is 0. If you set the performance buffer value to 0, the built-in algorithm is used. If you do not specify the performance buffer value, 0 is automatically used.

The default performance buffer value is 0. If you use the **default** option, the built-in algorithm is used. If you do not specify this command, the **default** option is automatically used.

Buffer Pools

In the architecture of Generation 2, Generation 3, and Generation 4 modules, receive buffers shared by a set of ports are called *buffer groups*. The receive buffer groups are organized into *global* and *local* buffer pools.

The receive buffers allocated from the global buffer pool to be shared by a port group are called a *global receive buffer pool*. Global receive buffer pools include the following buffer groups:

- Reserved internal buffers
- Allocated BB_credit buffers for each Fibre Channel interface (user configured or assigned by default)
- Common unallocated buffer pool for BB_credits, if any, to be used for additional BB_credits as needed
- Performance buffers (only used on 12-port 4-Gbps and 4-port 10-Gbps switching modules)



Note The 48-port and 24-port 8-Gbps modules have *dual global buffer pools*. Each buffer pool in the 48-port modules support 24 ports and in the 24-port modules each buffer pool supports 12 ports.

The following figure shows the allocation of BB_credit buffers on line cards (24-port and 48-port 4-Gbps line cards).

Figure 60-1 Receive Buffers for Fibre Channel Ports in a Global Buffer Pool

The following figure shows the default BB_credit buffer allocation model for 48-port 8-Gbps switching modules. The minimum BB_credits required to bring up a port is two buffers.

Figure 60-2 BB_Credit Buffer Allocation in 48-Port 8-Gbps Switching Modules

The following figure shows the default BB_credit buffer allocation model for 24-port 8-Gbps switching modules. The minimum BB_credits required to bring up a port is two buffers.

Figure 60-3 BB_Credit Buffer Allocation in 24-Port 8-Gbps Switching Modules

The following figure shows the default BB_credit buffer allocation model for 4/44-port 8-Gbps host-optimized switching modules. The minimum BB_credits required to bring up a port is two buffers.

Figure 60-4 BB_Credit Buffer Allocation in 4/44-Port 8-Gbps Switching Modules

The following figure shows the default BB_credit buffer allocation model for 24-port 4-Gbps switching modules. The minimum BB_credits required to bring up a port is two buffers.

Figure 60-5 BB_Credit Buffer Allocation in 24-Port 4-Gbps Switching Modules



Note The default BB_credit buffer allocation is the same for all port speeds.

BB_Credit Buffers for Switching Modules

This section describes how buffer credits are allocated to Cisco MDS 9000 switching modules, and includes the following topics:

Configuring Buffer Credits on a Generation 2, Generation 3 or Generation 4 Module

When you configure port mode to auto or E on a Generation 2 module, one of the ports will not come up for the following configuration:

- Port Mode: auto or E for all of the ports
- Rate Mode: dedicated
- Buffer Credits: default value

When you configure port mode to auto or E on a Generation 3 module, one or two of the ports will not come up for the following configuration:

- Port Mode: auto or E for the first half of the ports, the second half of the ports or for all of the ports
- Rate Mode: dedicated
- Buffer Credits: default value

When you configure port mode to auto or E for all ports in the global buffer pool, you need to reconfigure buffer credits on one or more of the ports. The total number of buffer credits configured for all the ports in the global buffer pool should be reduced by 64.

48-Port 8-Gbps Advanced Fibre Channel Module BB_Credit Buffers

The following table lists the BB_credit buffer allocation for the 48-port 8-Gbps Advanced Fibre Channel switching module.

Table 148: 48-Port 8-Gbps Advanced Switching Module BB_Credit Buffer Allocation

BB_Credit Buffer Allocation	BB_Credit Buffers Per Port		
	ISL	Fx Port	Fx Port
	Dedicated Rate Mode 8-Gbps Speed	Shared Rate Mode 8-Gbps Speed	
Default BB_credit buffers	250 for 48 port 500 for 32 port	32	32
Maximum BB_credit buffers	500	500	32

The following guidelines apply to BB_credit buffers on 32/48-port Advanced 8-Gbps Fibre Channel switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 500 buffers for dedicated rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 500 buffers for dedicated rate mode or 32 buffers for shared rate mode.
- Performance buffers are not supported on this module.
- The buffers should not be allocated automatically.

Each port group on the 32/48-port Advanced 8-Gbps Fibre Channel switching module consists of four/six ports. The ports in shared rate mode in a port group can have a maximum bandwidth oversubscription of 1.5:1 considering that each port group has 32-Gbps bandwidth. In case of 32 Port version, each port group of 4 ports has sufficient bandwidth (32 Gbps) to handle the line rate traffic without any oversubscription.

The following example configurations are supported by the 48-port Advanced 8-Gbps Fibre Channel switching modules:

- Six ports with shared rate mode and 8-Gbps speed (1.5:1 oversubscription) (default).
- Two port with dedicated rate mode and 8-Gbps speed plus four ports with shared rate mode and 8-Gbps speed (2:1 oversubscription).
- Two ports with dedicated rate mode and 8-Gbps speed plus four ports with shared rate mode and 8-Gbps speed (2:1 oversubscription) .
- One port with dedicated rate mode and 8-Gbps speed plus three ports with dedicated rate mode and 4-Gbps speed plus two ports with shared rate mode and 8-Gbps speed (1.33:1 oversubscription).
- Six ports with dedicated rate mode and 8-Gbps speed.

48-Port 8-Gbps Fibre Channel Module BB_Credit Buffers

The following table lists the BB_credit buffer allocation for the 48-port 8-Gbps Fibre Channel switching module.

Table 149: 48-Port 8-Gbps Switching Module BB_Credit Buffer Allocation

BB_Credit Buffer Allocation	BB_Credit Buffers Per Port		
	ISL	Fx Port	Fx Port
	Dedicated Rate Mode 8-Gbps Speed	Shared Rate Mode 8-Gbps Speed	
Default BB_credit buffers	250	32	32

BB_Credit Buffer Allocation	BB_Credit Buffers Per Port		
	ISL	Fx Port	Fx Port
	Dedicated Rate Mode 8-Gbps Speed		Shared Rate Mode 8-Gbps Speed
Maximum BB_credit buffers	500	500	32
Total Number of BB_Credit Buffers per Module			
Ports 1 through 24	6000		
Ports 25 through 48	6000		

The following guidelines apply to BB_credit buffers on 48-port 8-Gbps Fibre Channel switching modules:

- BB_credit buffers allocated for ports 1 through 24 and 25 through 48 can be a maximum of 6000 each so that the load is distributed.
- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 500 buffers for dedicated rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 500 buffers for dedicated rate mode or 32 buffers for shared rate mode.
- Performance buffers are not supported on this module.
- The buffers should not be allocated automatically.

Each port group on the 48-port 8-Gbps Fibre Channel switching module consists of six ports. The ports in shared rate mode in a port group can have a maximum bandwidth oversubscription of 10:1 considering that each port group has 12.8-Gbps bandwidth.

The following example configurations are supported by the 48-port 8-Gbps Fibre Channel switching modules:

- Six ports with shared rate mode and 8-Gbps speed (4:1 oversubscription) (default)
- One port with dedicated rate mode and 8-Gbps speed plus
five ports with shared rate mode and 8-Gbps speed (10:1 oversubscription)
- Two ports with dedicated rate mode and 4-Gbps speed plus
four ports with shared rate mode and 4-Gbps speed (4:1 oversubscription)
- One port with dedicated rate mode and 4-Gbps speed plus
three ports with dedicated rate mode and 2-Gbps speed plus
two ports with shared rate mode and 4-Gbps speed (4:1 oversubscription)
- Six ports with dedicated rate mode and 2-Gbps speed

24-Port 8-Gbps Fibre Channel Module BB_Credit Buffers

The following table lists the BB_credit buffer allocation for the 24-port 8-Gbps Fibre Channel switching module.

Table 150: 24-Port 8-Gbps Switching Module BB_Credit Buffer Allocation

BB_Credit Buffer Allocation	BB_Credit Buffers Per Port		
	ISL	Fx Port	Fx Port
	Dedicated Rate Mode 8-Gbps Speed	Shared Rate Mode 8-Gbps Speed	
Default BB_credit buffers	500	32	32
Maximum BB_credit buffers	500 ¹	500 ³⁶	32
Total Number of BB_Credit Buffers per Module			
Ports 1 through 12	6000		
Ports 13 through 24	6000		

³⁶ When connected to Generation 1 modules, reduce the maximum BB_credit allocation to 250.

The following guidelines apply to BB_credit buffers on 24-port 8-Gbps Fibre Channel switching modules:

- BB_credit buffers allocated for ports 1 through 12 and 13 through 24 can be a maximum of 6000 each so that the load is distributed.
- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 500 buffers for dedicated rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 500 buffers for dedicated rate mode or 32 buffers for shared rate mode.
- Performance buffers are not supported on this module.

Each port group on the 24-port 8-Gbps Fibre Channel switching module consists of three ports. The ports in shared rate mode in a port group can have a maximum bandwidth oversubscription of 10:1 considering that each port group has 12.8-Gbps bandwidth.

The following example configurations are supported by the 24-port 8-Gbps Fibre Channel switching modules:

- Three ports with shared rate mode and 8-Gbps speed (2:1 oversubscription) (default)
- One port with dedicated rate mode and 8-Gbps speed plus
two ports with shared rate mode and 8-Gbps speed (4:1 oversubscription)
- One port with dedicated rate mode and 8-Gbps speed plus
one port with dedicated rate mode and 4-Gbps speed plus
one port with shared rate mode and 8-Gbps speed (10:1 oversubscription)
- Two ports with dedicated rate mode and 4-Gbps speed plus
one port with shared rate mode and 8-Gbps speed (2:1 oversubscription)
- Three ports with dedicated rate mode and 4-Gbps speed

4/44-Port 8-Gbps Host-Optimized Fibre Channel Module BB_Credit Buffers

The following table lists the BB_credit buffer allocation for the 4/44-port 8-Gbps Fibre Channel switching module.

Table 151: 4/44-Port 8-Gbps Switching Module BB_Credit Buffer Allocation

BB_Credit Buffer Allocation	BB_Credit Buffers Per Port		
	ISL	Fx Port	Fx Port
	Dedicated Rate Mode 8-Gbps Speed	Shared Rate Mode 8-Gbps Speed	
Default BB_credit buffers	125	32	32
Maximum BB_credit buffers	250	250	32
Total number of BB_credit buffers per module	6000		

The following guidelines apply to BB_credit buffers on 4/44-port 8-Gbps Fibre Channel switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 500 buffers for dedicated rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 250 buffers for dedicated rate mode or 32 buffers for shared rate mode.
- Performance buffers are not supported on this module.

Each port group on the 24-port 8-Gbps Fibre Channel switching module consists of 12 ports. The ports in shared rate mode in a port group can have a maximum bandwidth oversubscription of 10:1 considering that each port group has 12.8-Gbps bandwidth.

The following example configurations are supported by the 4/44-port 8-Gbps Fibre Channel switching modules:

- Twelve ports with shared rate mode and 4-Gbps speed (5:1 oversubscription) (default)
- One port with dedicated rate mode and 8-Gbps speed plus
eleven ports with shared rate mode and 4-Gbps speed (10:1 oversubscription)
- One port with dedicated rate mode and 4-Gbps speed plus
three ports with dedicated rate mode and 3-Gbps speed plus
eight ports with shared rate mode and 4-Gbps speed (2:1 oversubscription)
- Twelve ports with dedicated rate mode and 1-Gbps speed

48-Port 4-Gbps Fibre Channel Module BB_Credit Buffers

The following table lists the BB_credit buffer allocation for the 48-port 8-Gbps Fibre Channel switching module.

Table 152: 48-Port 4-Gbps Switching Module BB_Credit Buffer Allocation

BB_Credit Buffer Allocation	BB_Credit Buffers Per Port		
	ISL ³⁷	Fx Port	Fx Port
	Dedicated Rate Mode 4-Gbps Speed	Shared Rate Mode 4-Gbps Speed	
Default BB_credit buffers	125	16	16
Maximum BB_credit buffers	250	250	16

BB_Credit Buffer Allocation	BB_Credit Buffers Per Port		
	ISL ³⁷	Fx Port	Fx Port
	Dedicated Rate Mode 4-Gbps Speed		Shared Rate Mode 4-Gbps Speed
Total number of BB_credit buffers per module	6000		

³⁷ ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 48-port 4-Gbps Fibre Channel switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- Performance buffers are not supported on this module.

Each port group on the 48-port 4-Gbps Fibre Channel switching module consists of 12 ports. The ports in shared rate mode have bandwidth oversubscription of 2:1 by default. However, some configurations of the shared ports in a port group can have maximum bandwidth oversubscription of 4:1 (considering that each port group has 12.8-Gbps bandwidth).

The following example configurations are supported by the 48-port 4-Gbps Fibre Channel switching modules:

- Twelve ports with shared rate mode and 4-Gbps speed (4:1 oversubscription) (default)
- One port with dedicated rate mode and 4-Gbps speed plus
11 ports with shared rate mode and 4-Gbps speed (5:1 oversubscription)
- One port with dedicated rate mode and 4-Gbps speed plus
11 ports with shared rate mode and 2-Gbps speed (2.5:1 oversubscription)
- Two ports with dedicated rate mode and 2-Gbps speed plus
10 ports with shared rate mode and 4-Gbps speed (5:1 oversubscription)
- Two ports with dedicated rate mode and 2-Gbps speed plus
10 ports with shared rate mode and 2-Gbps speed (2.5:1 oversubscription)
- Twelve ports with dedicated rate mode and 1-Gbps speed
- Three ports with dedicated rate mode and 4-Gbps speed plus
four ports with shared rate mode and 1-Gbps speed plus
five ports put out-of-service (see figure below)
insert image 144858.jpg
- Six ports with dedicated rate mode and 2-Gbps speed plus four ports with shared rate mode and 1-Gbps speed plus two ports put out-of-service (see below figure)
insert image 144859.jpg

24-Port 4-Gbps Fibre Channel Module BB_Credit Buffers

The following table lists the BB_credit buffer allocation for 24-port 4-Gbps Fibre Channel switching modules.

Table 153: 24-Port 4-Gbps Switching Module BB_Credit Buffer Allocation

BB_Credit Buffer Allocation	BB_Credit Buffers Per Port		
	ISL ³⁸	Fx Port	Fx Port
	Dedicated Rate Mode 4-Gbps Speed	Shared Rate Mode 4-Gbps Speed	
Default BB_credit buffers	250	16	16
Maximum BB_credit buffers	250	250	16
Total number of BB_credits buffers per module	6000		

³⁸ ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 24-port 4-Gbps Fibre Channel switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- Performance buffers are not supported on this module.

Each port group on the 24-port 4-Gbps Fibre Channel switching module consists of six ports. The ports in shared rate mode have a bandwidth oversubscription of 2:1 by default. However, some configurations of the shared ports in a port group can have a maximum bandwidth oversubscription of 4:1 (considering that each port group has 12.8-Gbps bandwidth).

The following example configurations are supported by the 24-port 4-Gbps Fibre Channel switching modules:

- Six ports with shared rate mode and 4-Gbps speed (2:1 oversubscription) (default)
 - Two ports with dedicated rate mode and 4-Gbps speed plus
four ports with shared rate mode and 4-Gbps speed (with 4:1 oversubscription)
 - One port with dedicated rate mode and 4-Gbps speed plus
three ports with dedicated rate mode and 2-Gbps speed plus
two ports with shared rate mode and 4-Gbps speed (4:1 oversubscription)
 - Six ports with dedicated rate mode and 2-Gbps speed
 - Three ports with dedicated rate mode and 4-Gbps speed plus
three ports with shared rate mode and 1-Gbps speed (see below figure)
- insert image 144857.jpg

18-Port Fibre Channel/4-Port Gigabit Ethernet Multiservice Module BB_Credit Buffers

The following table lists the BB_credit buffer allocation for 18-port 4-Gbps multiservice modules.

Table 154: 18-Port 4-Gbps Multiservice Module BB_Credit Buffer Allocation

BB_Credit Buffer Allocation	BB_Credit Buffers Per Port			
	ISL ³⁹	Fx Port	ISL	Fx Port
	Dedicated Rate Mode 4-Gbps Speed		Shared Rate Mode 4-Gbps Speed	
Default BB_credit buffers	250	16	16	16
Maximum BB_credit buffers	250	250	16	16
Total number of BB_credit buffers per module	4509			

³⁹ ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 18-port 4-Gbps Fibre Channel switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- Performance buffers are not supported on this module.

12-Port 4-Gbps Switching Module BB_Credit Buffers

The following table lists the BB_credit buffer allocation for 12-port 4-Gbps switching modules.

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port	
	ISL ⁴⁰	Fx Port
	Dedicated Rate Mode 4-Gbps Speed	
Default BB_credit buffers	250	16
Maximum BB_credit buffers	250	16
Default Performance buffers	145	12
Total number of BB_credit buffers per module	5488	
Total number of performance buffers per module	512 (shared)	

⁴⁰ ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 12-port 4-Gbps switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers.
- BB_credit buffers for Fx port mode connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers.
- By default, 512 performance buffers are preallocated and are shared by all the ports. These buffers are configurable and the buffers are assigned to the port based on the availability of the buffers in the shared pool.
- There are 2488 extra buffers available as extended BB_credit buffers after allocating all the default BB_credit buffers for all the ports in ISL mode (5488 - (250 * 12)).



Note Extended BB_credits are allocated across all ports on the switch. That is, they are not allocated by port group.



Note By default, the ports in the 12-port 4-Gbps switching modules come up in 4-Gbps dedicated rate mode but can be configured as 1-Gbps and 2-Gbps dedicated rate mode. Shared mode is not supported.

4-Port 10-Gbps Switching Module BB_Credit Buffers

The following table lists the BB_credit buffer allocation for 4-port 10-Gbps switching modules.

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port	
	ISL ⁴¹	F port ⁴²
	Dedicated Rate Mode 10-Gbps Speed	
Default BB_credit buffers	250	16
Maximum BB_credit buffers	750	16
Maximum BB_credit buffers on one of the ports with Enterprise license	4095	
Total number of BB_credit buffers per module	5488	
Default Performance buffers	145	12
Total number of performance buffers per module	512 (shared)	

⁴¹ ISL = E port or TE port.

⁴² Ports on the 4-port 10-Gbps cannot operate in FL port mode.



Note The ports in the 4-port 10-Gbps switching module only support 10-Gbps dedicated rate mode. FL port mode and shared rate mode are not supported.

The following considerations apply to BB_credit buffers on 4-port 10-Gbps switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 750 buffers.
- BB_credit buffers for Fx port mode connections can be configured from a minimum of 2 buffers to a maximum of 750 buffers.
- By default, 512 performance buffers are preallocated and are shared by all the ports. These buffers are configurable and the buffers are assigned to the port based on the availability of the buffers in the shared pool.
- There are 2488 extra buffers available as extended BB_credits after allocating all the default BB_credit buffers for all the ports in ISL mode (5488 - (750 * 4)).



Note Extended BB_credits are allocated across all ports on the switch. That is, they are not allocated by port group.

BB_Credit Buffers for Fabric Switches

This section describes how buffer credits are allocated to Cisco MDS 9000 Fabric switches, and includes the following topics:

Cisco MDS 9148 Fabric Switch BB_Credit Buffers

The following table lists the BB_credit buffer allocation for 48-port 8-Gbps Fibre Channel switches.

Table 155: 48-Port 8-Gbps Fabric Switch BB_Credit Buffer Allocation

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port Group	BB_Credit Buffers Per Port	
		ISL ⁴³	Fx Port
Default BB_credit buffers	128	32	32
Maximum configurable BB_credit buffers on 8-Gbps mode	128	125	125

⁴³ ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 48-port 8-Gbps Fabric Switches:

- BB_credit buffers can be configured from a minimum of 1 buffer to a maximum of 32 buffers per port when the ports are in F or FL mode.
- BB_credit buffers can be configured from a minimum of 2 buffers to a maximum of 32 buffers per port when the ports are in E or TE mode.

Cisco MDS 9134 Fabric Switch BB_Credit Buffers

The following table lists the BB_credit buffer allocation for MDS 9134 Fabric Switches.

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port Group	BB_Credit Buffers Per Port	
		ISL ⁴⁴	Fx Port
Maximum user-configurable BB_credit buffers	64	61	61
Minimum user-configurable BB_credit buffers		2	1
Default BB_credit buffers on 10-Gbps mode	64	64	64
Default BB_credit buffers on 4-Gbps mode	64	16	16

⁴⁴ ISL = E port or TE port.

Cisco MDS 9124 Fabric Switch BB_Credit Buffers

The following table lists the BB_credit buffer allocation for MDS 9124 Fabric Switches.

Table 156: MDS 9124 Fabric Switch BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port Group	BB_Credit Buffers Per Port Defaults	
		ISL ⁴⁵	Fx Port
Maximum user-configurable BB_credit buffers	64	61	61
Minimum user-configurable BB_credit buffers		2	1
Default BB_credit buffers	64	16	16

⁴⁵ ISL = E port or TE port.

Cisco MDS 9222i Multiservice Modular Switch BB_Credit Buffers

The following table lists the BB_credit buffer allocation for 18-port 4-Gbps Multiservice Modular switches.

Table 157: 18-Port 4-Gbps Fabric Switch BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port Group	BB_Credit Buffers Per Port Defaults	
		ISL ⁴⁶	Fx Port
User-configurable BB_credit buffers	4509	250	16

⁴⁶ ISL = E port or TE port.

Extended BB_Credits

To facilitate BB_credits for long-haul links, the extended BB_credits feature allows you to configure the receive buffers above the maximum value on all Generation 2, Generation 3 and Generation 4 switching modules. When necessary, you can reduce the buffers on one port and assign them to another port, exceeding the default maximum. The minimum extended BB_credits per port is 256 and the maximum is 4095.



Note

Extended BB_credits are not supported on the Cisco MDS 9148 Fabric Switch, Cisco MDS 9134 Fabric Switch, Cisco MDS 9124 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

In general, you can configure any port in a port group to dedicated rate mode. To do this, you must first release the buffers from the other ports before configuring larger extended BB_credits for a port.



Note

The ENTERPRISE_PKG license is required to use extended BB_credits on Generation 2, Generation 3 and Generation 4 switching modules. Also, extended BB_credits are not supported by ports in shared rate mode. All ports on the Generation 2 and Generation 3 switching modules support extended BB_credits. There are no limitations for how many extended BB_credits you can assign to a port (except for the maximum and minimum limits). If necessary, you can take interfaces out of service to make more extended BB_credits available to other ports.

You can use the extended BB_credits flow control mechanism in addition to BB_credits for long-haul links.

This section includes the following topics:

Extended BB_credits on Generation 1 Switching Modules

The BB_credits feature allows you to configure up to 255 receive buffers on Generation 1 switching modules. To facilitate BB_credits for long haul links, you can configure up to 3,500 receive BB_credits on a Fibre Channel port on a Generation 1 switching module.

To use this feature on Generation 1 switching modules, you must meet the following requirements:

- Obtain the ENTERPRISE_PKG license. See the *Cisco MDS 9000 Family NX-OS Licensing Guide*.
- Configure this feature in any port of the full-rate 4-port group in either the Cisco MDS 9216i Switch or in the MPS-14/2 module (see [Figure 187: Port Group Support for the Extended BB_Credits Feature, on page 1120](#)).

Figure 187: Port Group Support for the Extended BB_Credits Feature

insert image 120479 here

The port groups that support extended credit configurations are as follows:

- Any one port in ports 1 to 4 (identified as Group 1).
- Any one port in ports 5 to 8 (identified as Group 2).
- Any one port in ports 9 to 12 (identified as Group 3).



Note

The last two Fibre Channel ports (port 13 and port 14) and the two Gigabit Ethernet ports do not support the extended BB_credits feature.

- Explicitly enable this feature in the required Cisco MDS switch.
- Disable the remaining three ports in the 4-port group if you need to assign more than 2,400 BB_credits to the first port in the port group.
 - If you assign less than 2,400 extended BB_credits to any one port in a port group, the remaining three ports in that port group can retain up to 255 BB_credits based on the port mode.



Note

The receive BB_credit value for the remaining three ports depends on the port mode. The default value is 16 for the Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required without exceeding the maximum value of 255 BB_credits.

- If you assign more than 2,400 (up to a maximum of 3,500) extended BB_credits to the port in a port group, you must disable the other three ports.
- If you change the BB_credit value the port is disabled, and then reenabled.
 - Disable (explicitly) this feature if you need to nondisruptively downgrade to Cisco SAN-OS Release 1.3 or earlier. When you disable this feature, the existing extended BB_credit configuration is completely erased.



Note The extended BB_credit configuration takes precedence over the receive BB_credit and performance buffer configurations.

Extended BB_credits on Generation 2 and Generation 3 Switching Modules

To use this feature on Generation 2 or Generation 3 switching modules, you must meet the following requirements:

- Display the interface configuration in the Information pane.
- Obtain the Enterprise package (ENTERPRISE_PKG) license (see the *NX-OS Family Licensing Guide*).
- Configure this feature in any port on a Generation 2 switch module. See the [Extended BB_Credits, on page 1119](#) for more information on extended BB_credits on Generation 2 switching modules.



Note Extended BB_credits are not supported on the Cisco MDS 9124 Fabric Switch, Cisco MDS 9134 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

Buffer-to-Buffer Credit Recovery

Although the Fibre Channel standards require low bit error rates, bit errors do occur. Over time, the corruption of receiver-ready messages, known as R_RDY primitives, can lead to a loss of credits, which can eventually cause a link to stop transmitting in one direction. The Fibre Channel standards provide a feature for two attached ports to detect and correct this situation. This feature is called buffer-to-buffer credit recovery.

Buffer-to-buffer credit recovery functions as follows: the sender and the receiver agree to send checkpoint primitives to each other, starting from the time that the link comes up. The sender sends a checkpoint every time it has sent the specified number of frames, and the receiver sends a checkpoint every time it has sent the specified number of R_RDY primitives. If the receiver detects lost credits, it can retransmit them and restore the credit count on the sender.

The buffer-to-buffer credit recovery feature can be used on any non arbitrated loop link. This feature is most useful on unreliable links, such as MANs or WANs, but can also help on shorter, high-loss links, such as a link with a faulty fiber connection.



Note The buffer-to-buffer credit recovery feature is not compatible with the distance extension (DE) feature, also known as buffer-to-buffer credit spoofing. If you use intermediate optical equipment, such as DWDM transceivers or Fibre Channel bridges, on ISLs between switches that use DE, then buffer-to-buffer credit recovery on both sides of the ISL needs to be disabled.

Buffer-to-Buffer State Change Number

The BB_SC_N field (word 1, bits 15-12) specifies the buffer-to-buffer state change (BB_SC) number. The BB_SC_N field indicates that the sender of the port login (PLOGI), fabric login (FLOGI), or ISLs (E or TE ports) frame is requesting $2^{\text{SC_BB_N}}$ number of frames to be sent between two consecutive BB_SC send primitives, and twice the number of R_RDY primitives to be sent between two consecutive BB_SC receive primitives.

For Generation 2 and Generation 3 modules, the BB_SCN on ISLs (E or TE ports) is enabled by default. This can fail the ISLs if used with optical equipment using distance extension (DE), also known as buffer-to-buffer credit spoofing.

On a Generation-2 module, one port will not come up for the following configuration for all ports:

- Port Mode: auto or E for all the ports
- Rate Mode: dedicated
- Buffer Credits: default value

On a Generation-3 module, one or two ports will not come up for the following configuration for the first half of the ports, the second half of the ports, or all ports:

- Port Mode: auto or E for the first half of the ports, the second half of the ports, or for all of the ports
- Rate Mode: dedicated
- Buffer Credits: default value

When you configure port mode to auto or E and rate-mode to dedicated for all ports in the global buffer pool, you need to reconfigure buffer credits on one or more ports (other than default).


Note

If you use distance extension (buffer-to-buffer credit spoofing) on ISLs between switches, the BB_SCN parameter on both sides of the ISL needs to be disabled.

Receive Data Field Size

You can also configure the receive data field size for Fibre Channel interfaces. If the default data field size is 2112 bytes, the frame length will be 2148 bytes.

Configuring Interface Buffers

This section includes the following topics:

Configuring Buffer-to-Buffer Credits

To configure BB_credits for a Fibre Channel interface using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand **Switches > Interfaces**, and then select FC Physical. You see the interface configuration in the Information pane.
- Step 2** Click the **Bb Credit** tab.
You see the buffer credits.
- Step 3** Set any of the buffer-to-buffer credits for an interface.
- Step 4** Click **Apply Changes**.

Configuring Performance Buffers

To configure performance buffers for a Fibre Channel interface using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand **Switches > Interfaces**, and then select **FC Physical**.
You see the interface configuration in the Information pane.
- Step 2** Click the **BB Credit** tab.
You see performance buffer information in the Perf Bufs Admin and Perf Bufs Oper columns.
- Step 3** Set the performance buffers for an interface.
- Step 4** Click **Apply Changes**.
-

Configuring Extended BB_credits

To configure extended BB_credits for an MDS-14/2 interface, for a Generation 2 switching module interface, or for an interface in a Cisco MDS 9216i switch using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand **Switches > Interfaces**, and then select **FC Physical**. You see the interface configuration in the Information pane.
- Step 2** Click the **BB Credit** tab.
- Step 3** In the **Extended** column, set the extended BB_credits for the selected interface.
- Step 4** Click **Apply Changes**.
-

Configuring Receive Data Field Size

To configure the receive data field size using DCNM-SAN, follow these steps:

Procedure

- Step 1** Expand **Switches > FC Interfaces**, and then select **FC Physical**.
You see the interface configuration in the Information pane.
- Step 2** Click the **Other** tab and set the RxDataFieldSize field.
- Step 3** (Optional) Set other configuration parameters using the other tabs.
- Step 4** Click **Apply Changes**.
-



CHAPTER 61

Verifying Ethernet Interfaces

- [Verifying Ethernet Interfaces, on page 1125](#)

Verifying Ethernet Interfaces

This chapter includes the following sections:

Information About Ethernet Interfaces

DCNM-SAN 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 DCNM-SAN or Device Manager.

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

Default Settings

[Table 158: Default Ethernet Interface Parameters, on page 1125](#) lists the default settings for all physical Ethernet interfaces.

Table 158: Default Ethernet Interface Parameters

Parameters	Default
Oper Speed	10 GB
Admin Status	Up
CDP	True
VLAN Type	Static
VLAN List	1

Verifying Ethernet Interfaces Configuration

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

This section describes how to display the Ethernet interface status and includes the following topics:

Displaying Interface Information Using DCNM-SAN

To display Ethernet interfaces using DCNM-SAN, follow these steps:

Procedure

- Step 1** In the Physical Attributes pane, expand **Switches > Ethernet Interfaces > Physical**.
You see the Ethernet interface information pane.
The General tab displays the description, speed, MAC address, and status for each Ethernet interface.
- Step 2** Click the **VLAN** tab to display the VLAN assigned to each interface.
- Step 3** Click the **CDP Neighbors** tab to display the CDP neighbor assigned to each interface.
-

Displaying Interface Information Using Device Manager

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

Procedure

- Step 1** Launch Device Manager.
- Step 2** Choose **Interface > Ethernet > Physical**.
You see the Ethernet Interfaces information pane.
The General tab displays the description, speed, MAC address, and status for each interface.
- Step 3** Click the **VLAN** tab to display the VLAN assigned to each interface. Click the **CDP Neighbors** tab to display the CDP neighbor assigned to each interface.
-